Piloting of Multimedia Integrated Communications for European Researchers (MICE)

P.T. Kirstein
M.J. Handley
M.A. Sasse*

Abstract

The paper defines multimedia conferencing and describes the facilities currently available. It discusses briefly the activities which require standardisation, and the progress in this direction to date. It gives an overview of the MICE project, which utilises existing conferencing rooms, workstations, codecs and software, and existing network infrastructure, to offer researchers conferencing facilities within Europe, as well as a link to the US. The goals of the project, its achievements to date, and problems encountered are discussed in detail. Finally, we outline forthcoming activities.

I Introduction

In 1992, the CEC agreed to finance a one-year project called "Piloting of Multimedia Integrated Services for European Researchers (MICE)", which started on December 1, 1992. This paper will outline what we are trying to achieve, the technical choices we have made, and our progress up to June 1993.

Multimedia Conferencing, in the context of the activities to be discussed in this paper, implies that voice, video and computer data are to be shared amongst geographically distributed groups of people. Much research has been done in the individual areas; the communication service providers (PTOs) now offer video/voice facilities as commercial services, whilst communication equipment providers are offering video-conferencing studios complete with workstations. Hence it is reasonable to ask why the CEC should sponsor a Pilot Project in this area.

In practice, most of the facilities which are currently available are in one of the following classes:

1. High quality studios, often used by large commercial organisations, and connected by high quality codecs at 1.5 Mbps (T1), 2 Mbps (E1), or by analogue TV channels. The communications channels are circuit-switched, and usually only two-site point-to-point circuits, though sometimes multi-point is possible;
2. Simpler, lower quality studios connected over lower rate circuit-switched channels. These usually operate between 384 and 1024 Kbps, though more recently some use has been made of Basic Rate ISDN (64 or 128 Kbps) to connect such studios.
3. Economic workstations (just emerging) which operate using one or two Basic Rate ISDN channels (64 or 128 Kbps); if one channel only is used for the A/V, it is often possible to put data on the second channel of a Basic Rate ISDN connection. Such workstations usually incorporate a codec board to perform video compression.
4. Multimedia conferencing facilities have been set up on higher-end workstations. Software compression of video in such workstations is possible, typically at speeds of between 64 and 512 Kbps. Usually, the frame rate or picture quality is lower than video from hardware at comparable rates, due to limitations on available processing power. Such systems [1,2] have been deployed to a limited extent in the research communities on wide area networks. The systems are particularly interesting, as they operate over packet-switched networks rather than circuits.

In commercial services, one usually puts in dedicated leased lines to meet the communication needs, or the Public Telecommunications Operators (PTOs) allow rental of the facilities on demand. Since there is often some difficulty in leasing the sorts of facilities described in (III), these are often carved out of larger, leased, private facilities.

In conference rooms systems, the traffic is usually carried over circuit-switched channels. This requires a constant bit-rate service for video, and that any voice or data facilities have a negligible impact on the additional communication requirements. It has been shown [3] that this assumption is erroneous, and that there can be consider-

*The authors are with University College London. They may be reached at M.Handley@cs.ucl.ac.uk

able advantages in packet-switched operation — though this may be largely irrelevant as we move to higher speed, broadband ISDN, Asynchronous Transfer Mode (ATM) services.

Under DARPA auspices, there are several facilities for providing multimedia facilities over the Internet; DARTNET [3] and the Defence Simulation Internet are two examples. These operate with codecs, and different mechanisms for conference control and multiplexing have been incorporated. Typically they operate at 128 — 384 Kbps speeds, with bandwidth allocations in the routers (ST-II) to ensure that adequate performance is achieved so that the codecs do not lose synchronisation. The European research networks, whilst being more heterogeneous and generally of lower bandwidth, have much the same technological characteristics as the US part of the Internet; it would be most convenient if multimedia facilities could be provided for European researchers over this communication infrastructure preferably with links also to the US. The MICE project aims to provide just such facilities on a small scale.

In the European context, there are some significant differences to the older US activities. First there is a recent international standard for codecs (H.261 for video coding, H.221 for multimedia streams over serial lines, and various other standards for the audio coding); in accordance with European policies, we will use these standards [4] to the extent practicable. Secondly in addition to the packet research networks, ISDN [5] is now widely available; the speeds of 128 - 256 Kbps available on two - four such channels is approximately the same as speeds available over the international packet networks. We expect to provide access and coordination for conference rooms or workstations using H.261/H.221 codecs operating both over the research packet-switched networks and the ISDN.

Over the last few years, workstations have become much more powerful. Manufacturers have been determined to develop new applications, and have decided that multimedia has tremendous sales potential. Most of their activities have concentrated on local area networks (because the cost of codecs has been prohibitively expensive), but they have recognised also the need for interoperation and standardisation. Two standards have emerged: JPEG (the Joint Photographic Experts Group, [6]) for still pictures, and MPEG (the Motion Picture Experts Group [7]) for moving pictures. These provide very good results at speeds above about 512 Kbps, and again specialised chip sets are becoming available.

Even without the special chipsets, current workstations are capable of providing codec functionality in software - albeit with reduced performance. Such software is often aimed purely at packet network users. There are many users who want multimedia facilities and who are willing to tolerate low frame rate video which can be provided by software codecs. Some manufacturers are also starting to provide hardware assistance for compression, which further blurrs the differences between workstations with and without codecs.

In summary, the current situation requires facilities for users of both conference rooms and workstations, and connectivity between the two approaches. Much of early part of this paper has been presented in [8]; for completeness, the general overview material is reproduced here. That paper should be regarded, however, as complementary to this presentation.

The MICE project is funded at the modest level of 54 man-months during its first year. It would be quite impractical to undertake this ambitious project with that level of resources if it was not based on other very extensive activities. This is not the place to describe all such work, but some of the more relevant is referenced here. It is important to note that we are using both software developed by the MICE partners in other projects, and also some developed by colleagues elsewhere.

II The MICE Aims and Participants

The MICE project, coordinated by UCL, addresses these requirements; an early description is given in [8]. It aims to provide multi-way integrated conferencing services between a number of European pilot sites, and link them to the appropriate communities in the US. MICE is demonstrating a range of possible ways of interworking between the partners (who are using heterogeneous environments) rather than to undertake a large development activity. The intention is to interconnect up to half a dozen conference rooms, and up to a dozen workstations interconnected both by the ISDN and research networks, and to pilot multi-way conferencing services between them. The project aims to demonstrate point-to-point conferencing between 5 sites in 5 countries within six months, and multipoint with up to 4 simultaneous users from 12 sites in 7 countries within 12 months. The progress we have made to date is indicated in this paper; at the conference, we will give more current information on our progress.

The initial partners in the MICE project are:

Stuttgart Univ (Germany),	NTR (Norway),
Univ of Oslo (Norway),	SICS (Sweden),
ULB (Belgium),	INRIA(France),
GMD (Germany),	ONERA (France),
Nottingham Univ (UK),	UCL (UK).

In addition, two National Research Network Operators (the German DFN and the British JNT) are associated with the project without funding. Both they and the CEC are planning to attach their own conferencing systems for interactions with ones provided in MICE. Others have asked to join the project, and we will consider shortly how this can be achieved without jeopardising the attainment of its ambitious short-term goals.

The MICE project relies heavily on the developments, expertise and facilities provided in other national and international projects; only the international aspects, the technology integration, and the operational coordination are financed by the CEC. Therefore the project uses a range of existing conferencing equipment from other projects; this has a major bearing on the choice of subsystems used. Since so many countries and public monies are involved, we have chosen to adopt international standards where possible.

III Multimedia Conferencing and the Responsibilities of the Partners

III.A Overview

As stated in the previous sections, voice, video and computer data are to be shared amongst a distributed group of people. In the MICE project the users are part of the research community. We have divided the requirements into four classes of service:

1. Multi-way video
2. Multi-party Voice
3. Multi-way shared workspace
4. Multiplexing and Conference Management

This is not the only way that the services could have been divided - it is explained below why this split was chosen. Similarly, the existing research networks are acting as carrier networks - and these are packet-switched services. Details of this infrastructure are given in Section IV.

Full bandwidth video requires about 120 Mbps bandwidth; this would be quite unaffordable, and certainly could not be carried by the existing research networks. We have to use, therefore, compressed video. There are a large number of compression schemes available, but since the MICE project is committed using international standards where possible, we have agreed to use the international H.261 standard [4] as the core video compression scheme. This can be provided by special codecs, or at a reduced rate by software in some of the more powerful workstations. Both approaches will be pursued.

We do not expect to have enough bandwidth available to carry multi-way video to all sites, and even if we did, most hardware codecs can only decode one video stream at a time. For these reasons, we will use a hublet approach for at least some of the conferences; here all the video goes to one site, where it is multiplexed, and retransmitted to all the participants. The functions of the hublet are discussed below. One function is to relay between different coding schemes - e.g between stations using codec hardware and reduced performance ones doing the codec function by software.

The basic philosophy has been adopted that participants will be free to use any equipment provided it is conformant to certain external standards. As no standard for transmission of H.261 video over packet networks existed, the MICE partners have produced an internet draft standard [9], which is based on the internet draft standard RTP protocol [10]. However there will be a reference implementation which is documented and supported by one or more named participating organisations.

III.B Audio-visual Coding

The set of Standards called p x 64 [4] cover video, voice, data; they include the H.261 Standard for compressed video, several for voice, and H.221 for multimedia streams over serial links. Usually, data and voice are multiplexed within an H.221 stream - which is isochronous. The native H.261 coding for video is not strictly isochronous - though there are delay bounds. At present, most of the workstation-based software implementations treat H.261 video and audio separately. This allows us to route the video and audio separately; this feature is used in different ways. The audio, for instance, can be sent via direct internet multicast, while the video is sent to a Conference Multiplexing and Management Centre (CMMC); alternately, the video and audio can be given different priorities, so that sites on low bandwidth links only receive the audio.

A schematic of the architecture of the H.221/H.261 and related audio standards is shown in figure 1.

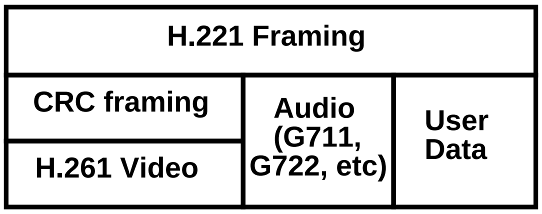

Figure 1: H.261 Video, Audio and User Data Combined in H.221 Framing

Because the H.221/H.261 stream has been designed for circuit switch environments, it is assumed that there is a capability to accept continuous streams of data; its framing is not very well suited for input to a workstation [11]. It is much more onerous to encode video than to decode it; while only a powerful workstation could be considered for a full H.261 encoding, decoding requires an order of magnitude less processing power. Where software implementations of H.261 exist, they are usually performed by omitting to perform a search for motion vectors; this results in a slightly lower quality video, but brings the processing requirements within the bounds set by today's workstation technology without special hardware assist. For example, on a 60MIPS workstation with a fast framegrabber, we can codec between 5 and 20 frames per second (depending on the amount of motion in the image) of QCIF video without motion compensation.

For a point-point conference over a circuit-switched network, it makes sense to integrate voice and video; for multi-way conferencing over a packet-switched net, this is less suitable. Voice is much more sensitive to delay and jitter than the video; moreover, it can be compressed to use much less bandwidth (between 4.8 - 16 Kbps depending on the required quality). By using silence suppression, and the fact that usually only one person is trying to talk at any time in a conference, it is possible to serve a large audience with voice for a modest outlay of bandwidth provided multicast is used. This has been demonstrated during all Internet Engineering Task Force (IETF) meetings since early 1992; several hundred sites in dozens of countries in five continents have participated - even though the same demonstrations have shown up problems in the networks and routers which can cause unacceptable delays and packet losses.

Often it is possible to engineer the packet switched networks to ensure reasonable bounds on delay and jitter in the research networks at our disposal [12]. We can also implement voice relays which mediate between sites with different coding schemes - a higher compression (lower quality) coding scheme is adopted to reduced communication bandwidths to some sites. While the most common method for encoding voice is PCM, there are several coding standards for reducing its bandwidth; examples are CELP [13], ADPCM [14], LPC [15], GSM [16]. We expect that many of these will be used; several are supported already in the audio tool we are using [17].

III.C Conference Rooms

A typical conference room will allow some 5 - 10 people to participate in a meeting. There are several versions of these in the participants' premises - with differing levels of sophistication. They normally include at least two remotely controllable cameras (one onto the participants and one vertically onto documents), a workstation (for camera control and for shared workspace activity), a number of microphones, several TV monitors or a TV projector, some small TVs for previewing, and a control box (or workstation software) to control all the peripherals. There is usually some device for packetising data from the cameras (usually a framegrabber in the workstation). To match the video speeds to those available to this project, there is normally a hardware codec.

A diagram of a typical CR setups is given in Fig. 2. Here the switching between cameras and displays is shown as affected by a video switch. This allows also other multimedia information to be introduced into the conferencing such as information from a video store. In the packet network case, the multimedia data coming into the CR workstation, which acts to strip the packet framing, pass the video to the codec and the decoded audio to the speaker, controls the video switch, and multiplexes and packetises in the other direction. This type of installation is reasonably expensive, so that its usage is normally shared between different functions.

The MICE organisations will be using their existing conference rooms. In addition, GMD has the responsibility of providing a reference installation which could be copied by organisations which do not have existing facilities but wish to join the programme.

III.D Workstations

Since it is possible to design a conference room around a workstation with projector, there is little intrinsic difference between a conference room and workstation (WS) installation. The workstation installation will normally be a cheaper installation. It is normally designed for one or two people, and has reduced functionality - and

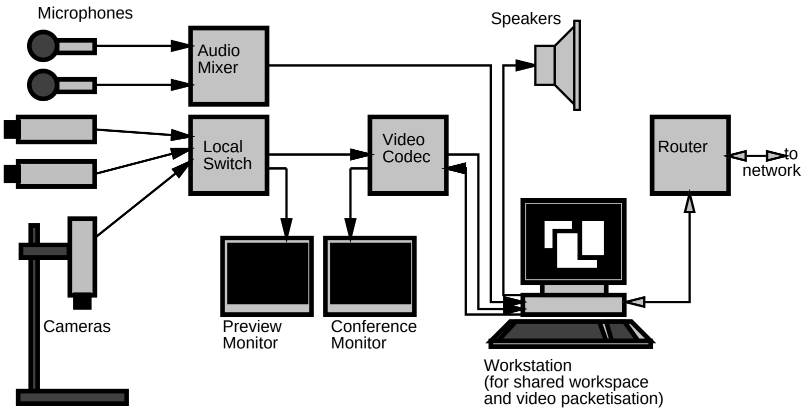

Schematic of Conference Room with Workstation providing media packetisation for transmission over a packet network
The audio compression and decompression is shown being performed in the workstation. Although the codec usually provides this functionality, compatibility with other sites is made simpler by using the workstation to provide audio functionanlity.

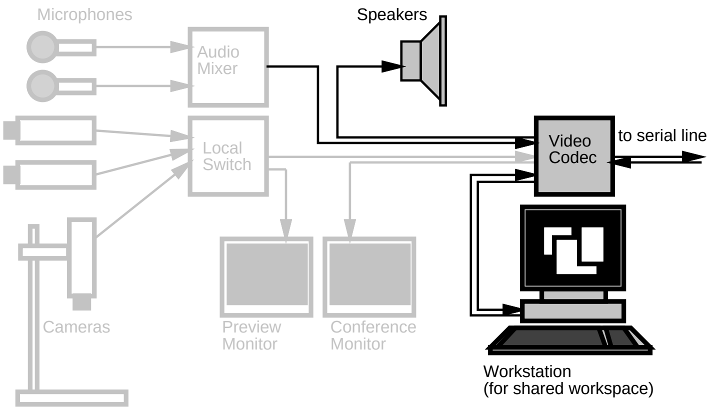

Figure 2: Schematic of Conference Room with Codec providing serial line framing for audio, video and shared workspace data.
(unchanged connections are shown in grey)

hence cost. There is often only one manually adjustable camera feeding a video frame-grabber, with compression usually being performed in software. Video display can be done on the screen of the workstation, though an external TV monitor will often be used if a hardware codec is chosen. The compression/decompression of the video in software is usually to a lower standard than that produced by hardware codecs. With the newer fast workstations, however, the quality is quite similar, at least at lower data rates (currently up to about 128 Kbps).

We also expect to accommodate workstations with reduced functionality which operate only with audio, or in which the video function is restricted to document entry and display. For this functionality, it may not even be necessary to include the camera/frame grabber, though some form of local scanner (e.g. a facsimile device interfaced to the workstation) would be helpful.

A diagram is given in fig. 5. The MICE organisations will be using their existing workstations; their audio software will be based on LBL's Visual Audio Tool (VAT) [17] and their video on the INRIA IVS system [1].

III.E Shared Workspace

There are many methods of supporting shared workspace activity. The partners are already supporting three different systems from other projects: CAR from UCL, Xspy from SICS, and Mscrawl from INRIA. The important activity here is that a broad range of application packages can be supported on the system. Some applications best support a large number of sites watching and listening to a conference; others best support a small number of groups interacting closely. Another aspect of shared working is the need for Conference Control. In MICE we are pursuing two approaches; a centralised and distributed approach. The centralised one will be implemented first by UCL and will form a part of the CMMC (see III.F). The current plan is to use an extended version of the CAR conference control system [18]. This will be used to interface

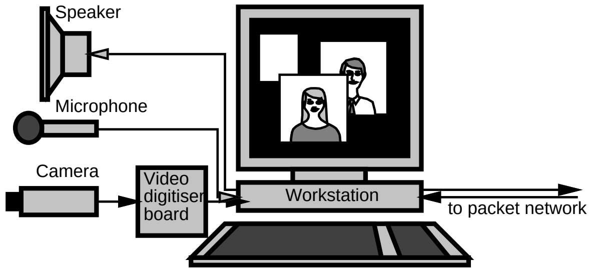

Multimedia Workstation performing the codec function in software

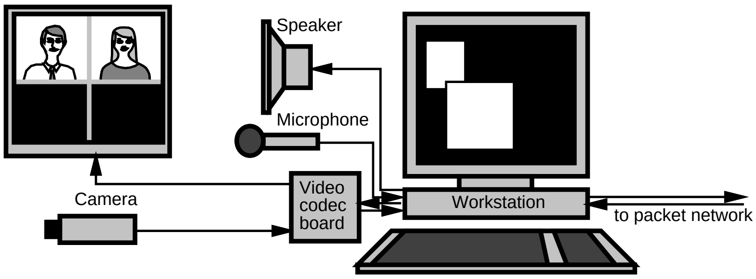

Figure 3: Multimedia Workstation performing the video codec function in hardware and the audio codec function in software
(instead of a separate monitor, a video overlay board can be used to overlay the video output from the codec board over a window on the workstation screen)

with a booking system for resource allocation, and will provide a directory service, conference setup, membership control, hublet resource allocation and basic floor control functionality where required. On top of this platform, various applications can be integrated. We intend to integrate in other shared workspace tools. This functionality will usually be added to the controlling workstation of the CR of Fig. 2 and the WS of Fig. 3.

III.F Conference Multiplexing and Management Centre

There will be at least one Conference Multiplexing and Management Centre (CMMC). Its functions include the following:

- The mechanism for allowing communities to come in from different networks (e.g. EMPB, UK-US Pipe, EBONE and ISDN);
- A way to relay between different protocol and coding standards;
- The incorporation of stations with reduced facilities (e.g. Basic Rate ISDN), video workstations without hardware codecs;
- A centre for multiplexing video and voice;
- A centre for conference reservation, resource control and management.
- A centre for ensuring multicast of the shared workspace data, and possibly the voice, to all conferees.

UCL has responsibility for the Conferencing Multiplexing and Management Centre (CMMC, UCL). A schematic representation is shown in Fig. 4. Here we show information as coming in over the ISDN or packet networks into one of two Routers (though the functions could be put into one router). The routers could be attached to the CMMC computing engine by a LAN - or could be even inside that engine. A more detailed description of the CMMC design is given in [8].

Where it is determined that incoming packets represent video to be decompressed by hardware, they can be passed through to a codec; where software preprocessing is required (e.g. because the information has been generated by incompatible software), this must be done prior to passing into the codec. At present, the video information from the US will be in a format incompatible to that from the European sites, and so must be passed to a particular codec. The function of the video multiplexor is to combine the video into quadrants; the compressed composite video is then returned to the server for multicast to all the conferees. Video from workstations with software codecs can also be passed also into the video multiplexor; on the return, the data for such workstations is may either be coded by a hardware codec in a low data rate mode, or alternatively by a software codec within the server. The connection with the UCL conference rooms can be direct via the video

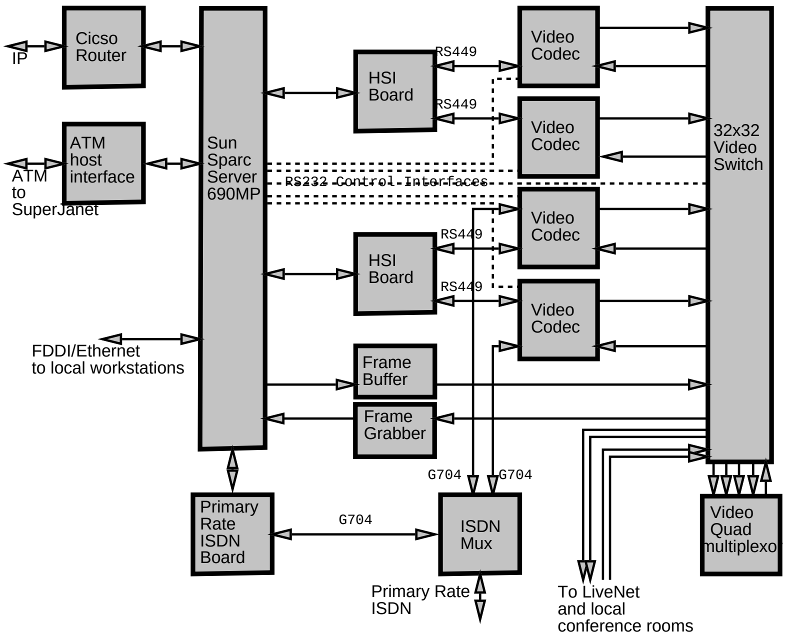

Figure 4: Schematic of the MICE Conference Management and Multiplexing Centre
(for simplicity, software codecs and hardware N-1 audio mixing are not shown)

switch.

With ISDN, the audio is usually sent in-band. In this case, the audio must be separated out for digital mixing - either in special hardware or in the computing engine. If only audio multicast directly from the participating sites is used, this functionality may not be needed. We do envisage, however, a need to relay between different audio coding systems - some providing better quality audio at higher bandwidths. In some cases we expect to use a combination of both centralised and decentralised mixing.

Many of the functions do not need to be centralised, and it is planned that the CMMC will eventually be replicated. However, the questions of which functions can be replicated, and how CMMCs should interact, require further study. One activity which will clearly be replicated is that of audio relay; we expect that there will be a deployment of facilities which can perform this function attached to, or as part of, many of the National networks.

IV The Infrastructure

No special communications facilities are being installed for MICE at first, although specific links are being upgraded by some of the participants. A diagram of the communications is shown in Fig. 5. There are currently three International networks involved: EuropaNet, EBONE and the Internet. In addition, the CMMC will support communications over ISDN. In each country there are National research networks, and these are connected to the International ones.

The international connections (with the National network in parentheses are:

- EBONE: France (Renater), Norway and Sweden (Nordunet), UK (JANET)
- EuropaNet: Belgium, Germany (DFN), UK (JANET)
- Internet → US: France (Renater), Sweden (Nordunet), UK (JANET)

The European network links are constantly being upgraded. The current status of the links is the following: EuropaNet links in the above are running at 2 Mbps in Germany, the Netherlands and the UK; the EBONE links are at about 1 Mbps between France and Scandinavia, but at significantly less to the UK. The Internet links run at 1 - 2 Mbps. There 30 64 Kbps channels of ISDN (Primary Rate) to UCL. However, it should be noted that the EBONE and Internet links also have substantial normal IP traffic; we expect that the EuropaNet links will be relatively lightly loaded during 1993. Direct links between EBONE and EuropaNet are scheduled to be complete shortly.

V The Status of Standards

We have already stated that we expect to use H.261 for the video coding. We know that the H.221 multimedia stream encoding is inappropriate for packet-switched operation, but to enable hardware codecs producing H.221 framed streams

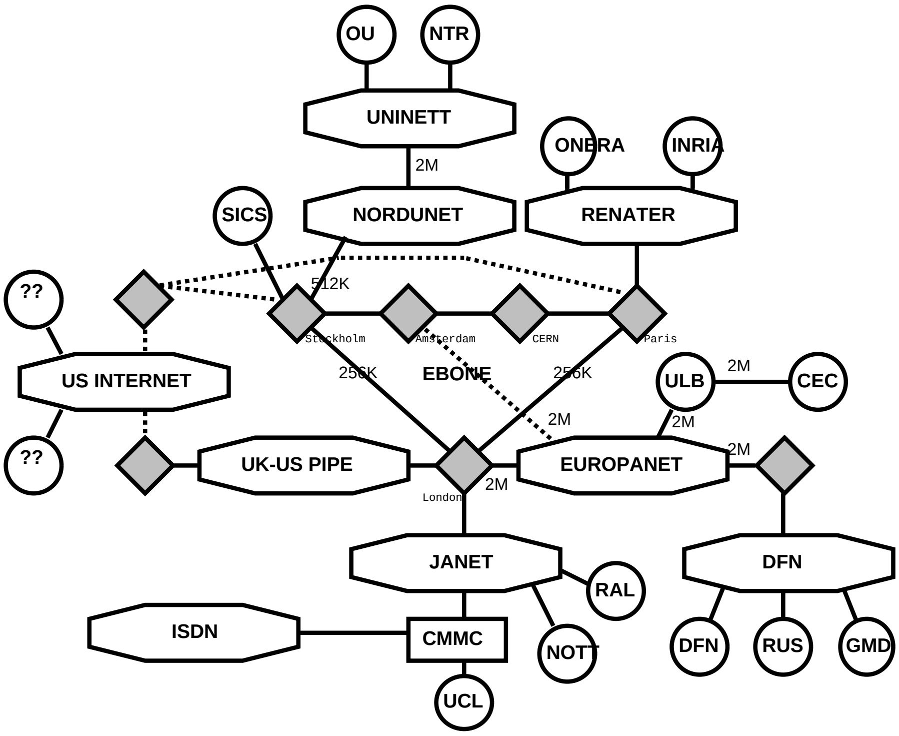

Figure 5:Schematic of International MICE Connectivity

to be used, we have developed workstation software to add and remove H.221 framing. We have also stated already that we doubt that the standards for in-band audio coding are appropriate, because of our potential need to use a different network topology for the voice than the video; some of the distribution sites require only the audio, others require reduced bandwidth audio with or without the video. Removing H.221 framing also satisfies this requirement by giving us a separate audio and video streams. We intend to use international standards again for the audio - but these may come from a different domain such as cellular telephone (GSM coding). There are emerging standards for Multipoint Control Units (MCUs), required for multiple conferences; these are supplementary to H.221/H.261. We will not be using them, as they are only appropriate if we are using H.221 framing. Some Internet Standards developed currently under the auspices of the Audio Visual Transport Working Group (AVT-WG) of the Internet Engineering Working Group (IETF) seem more appropriate. As there was no existing standard for transport of H.261 data over packet networks, we have created an internet draft standard [9]. This is fully conformant with the emerging RTP [10] internet standard for real time traffic. They also allow the specification of lifetimes for the elements, so that they can be destroyed if they are going to arrive too late at the destination - or alternately so that the priority of the elements can be changed at intermediate nodes. Of particular importance is that the coding permits different protocol versions to be used, which in turn enables heterogeneous implementations. Other features supported are flow control, end-to-end real-time control, media source and destination, quality of service, conference announcement, and quality of service announcement. It is significant that the working group includes representatives from supplier organisations like ATT.

Another set of standards are currently being developed by other IETF groups in Remote Conferencing (REM-WG) [19] and Conference Control. It is unlikely that these will be finalised for the early stages of MICE, but they may be incorporated later. There is already a concrete proposal that MICE be associated with a US NSF project, which would use the IETF Standards. The standards being developed in the working groups include the following: a Connection Conference Manager; an Applications Programming Interface; mechanisms for coding, transport, multicast and flow control; network resource management; multicast address administration; a connection and configuration manager, and conference systems management.

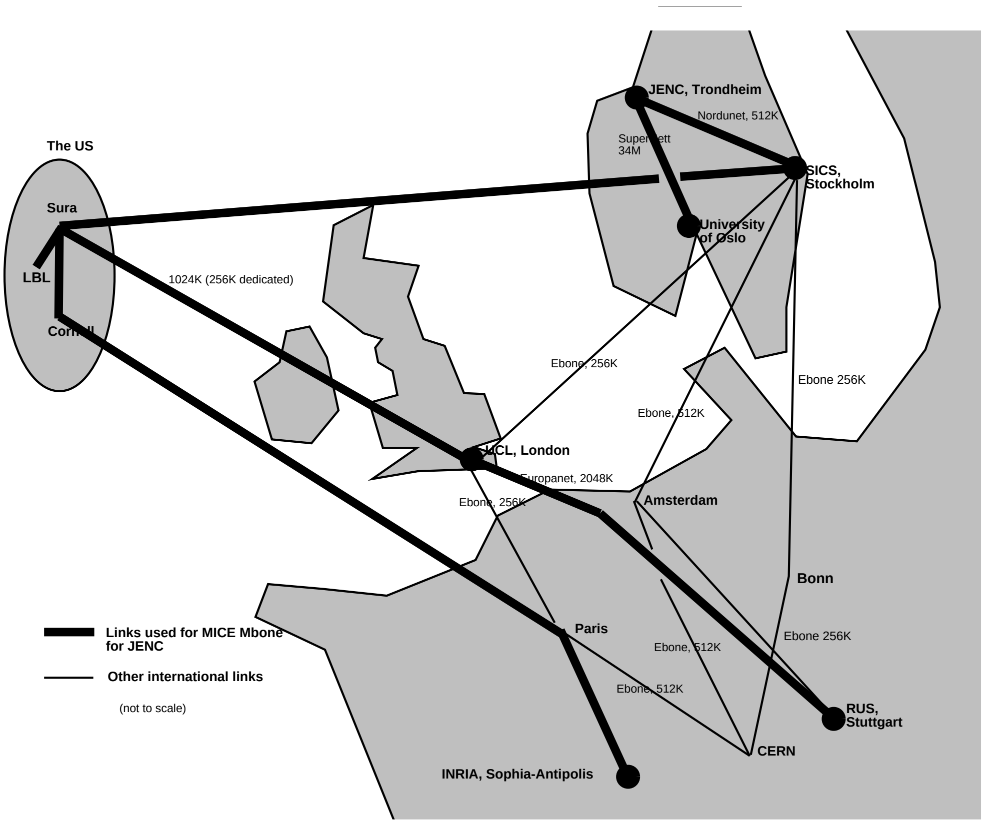

Figure 6: The Mbone connections used for the JENC demo

VI The Piloting Phase

VI.A Introduction

In the formulation of the project, we set various deliverables and milestones in the expectation that it would be important to show specific functionality. Moreover, the network infrastructure was specifically excluded from being an active subject of endeavour. In practice, it has been easy to demonstrate the sort of functionality proposed between specific sites at certain times; it has been much more difficult to obtain reliable service to all the sites - mainly because of the interplay between the stringent performance demanded by the applications and the inadequate performance of the network components.

In most European collaborative projects, even ones in the communications area, the facilities developed during the projects have not been used to further their completion; moreover, the demonstrations of their achievements have usually been elaborately put together for public displays. The main achievements have been paper deliverables and local demonstrations. We have decided on a different approach. We now use IVS [1] and VAT [17] as part of our regular weekly coordination meetings. We use the MESSIE [20] document system and its store to hold our project documents. We are targeting specific events to ensure that our level of capability takes a sustained step forward each time.

So far, we have concentrated on two such events. The first was at the Joint European Network Conference in Trondheim, Norway, May 10-13; the second at the Internet Engineering Task Force meeting in Amsterdam, the Netherlands in July 12-16. During the first we showed that the technology worked - but the different European network infrastructures were not yet working well together; at the second we are trying to ensure that the network infrastructure will really be adequate for the applications. We will target another event in the autumn; the main objectives of the autumn event will be influenced strongly by our experiences in July. We discuss below both what was shown in May, and what we are planning in July.

VI.B The JENC Demonstration

At JENC, we had the network infrastructure of Fig. 6. The main sites involved were the following (with the relevant networks in parentheses):

- INRIA (Sophia-Antipolis, France; Renater),
- LBL (Berkeley, US; Internet),

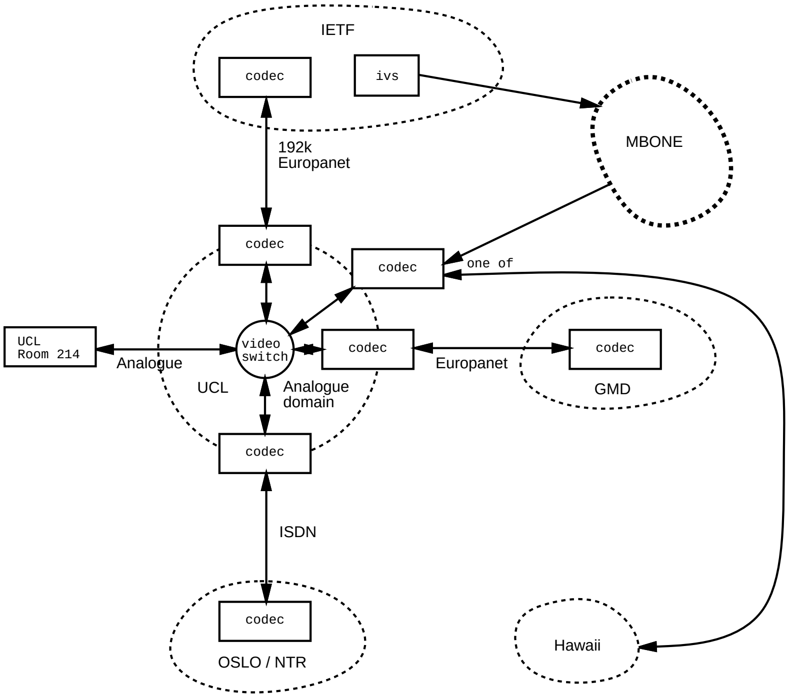

Figure 7: Schematic of the IETF demonstration

- JENC (Trondheim, Norway; Nordunet, SuperNett),
- University of Oslo (Oslo, Norway; SuperNett),
- RUS (Stuttgart, Germany; DFN),
- SICS (Stockholm, Sweden; Nordunet, EBONE),
- UCL (London, UK; SuperJanet).

In addition, the French, German and Swedish sites were attached to their international gateways (IG); the French and Swedish IGs supported EBONE and links to NSFNET; the British and German IGs supported EuropaNet; the British supported also a dedicated channel to the US DSI at 256 Kbps and a shared channel to NSFNET. Most of the international links were shared with other IP traffic; the Nordunet link was at 512 Kbps, the three NSF-European links were at speeds of about 1 Mbps. Some of the national segments were even at 34 Mbps (Norway and UK). There were many other international links - but these were the active ones. Although there is are direct EBONE links between the UK, France and Stockholm, these were too loaded to carry A/V traffic. All the traffic between Scandinavia and the rest of Europe had to cross the Atlantic twice.

At this event there were no hardware codecs; most of the sites had merely workstations running IVS [1] and VAT [17], though at Oslo and UCL there were conference rooms. Most of the workstations were using SparcStations with video cards (mainly SS-10s though some IPXs, mainly with Videopix cards though some Parallax); there were also a DECstation 5000 and a Silicon Graphics Indigo used for the multicast of the conference itself. Many intermediate points ran SparcStation or DECstation routers with mrouted [21] to provide multicast routing.

When all the equipment and links were working, good multiway working was achieved with all of the above sites (though usually less than six at a time). However, when there were problems with almost any of the links or routers, the service quickly deteriorated. For example, during breaks in the JENC conference when there was heavy data traffic on the Nordunet link, the packet loss became excessive. There were periodic downtimes in the US between Washington and FIX-East (where the UK link terminated); this resulted in INRIA, RUS and UCL becoming unreachable. There were some failures in the local loop in London - resulting in RUS and UCL becoming unreachable. We can say we demonstrated the multiway potential, but also the fragility of the system. Shared workspace and the real functionality of the CMMC were not demonstrated.

VI.C The IETF Demonstration

The IETF demonstration is intended to be much more ambitious than the JENC one. It will be sited at Amsterdam, which has access to both EuropaNet and EBONE - and the EBONE links will have been upgraded to 1.5 Mbps. The topology is shown in figure 7. We also expect to demonstrate some of the functionality of the CMMC, shared workspace, the use of hardware codecs, and interworking of hardware and software codecs. Part of the CMMC functionality that we expect to demonstrate is the interconnection of sites on the internet and a site connected over an international ISDN connection from Norway. We believe this will be the first time a single conference has demonstrated audio, video, and shared workspace interworking between both networks.

VII Conclusions

Multimedia conferencing is currently at a very exciting stage. There is a plethora of activities, and a large number of different approaches. There are an emerging set of standards - many truly appropriate, many starting to lead to convergence of approaches. The research community has a different environment, and different requirements, from the commercial services being developed by the carriers. The MICE project represents an important new initiative. It is based on the best practices of the research community. It will take maximum advantage of the networks provided by and for that community. It tries to set a framework that could be operated on a pan-European basis. It attempts to take advantage of the best practices in the US as well - and to provide a link between the European and North American activities. By its using its product to facilitate collaboration amongst its partners, it will ensure that its final product is usable. Since the project relies on the networks provided for the research community, it has to push the development of those networks until they meet the community's needs. MICE will impact equipment suppliers, service providers, communications carriers and the research community itself. This should be achieved through adopting the following measures;

- to use both specialised and generalised hardware;
- to incorporate both high end facilities but also reduced capabilities;
- to use both the research networks and the ISDN; to use both specialised conference rooms and ubiquitous workstations;
- to start in the more technologically advanced institutions, but to extend to less advanced institutions and less prepared countries with packaged solutions;
- to have researchers as users developing the facilities, but also those from other disciplines and from the funding bodies like JNT, DFN and the CEC itself.
- to start deploying more distributed relay facilities for all the relevant services.

The project is extremely challenging - and at the same time highly rewarding. Since it crosses many institutional and national boundaries, it may well create a precedent in removing such boundaries in the individual countries, on both sides of the Atlantic.

VIII References

[1] T. Turletti, "H.261 Software Codec for Videoconferencing over the Internet," *Research Report no 1834,* INRIA, Sophia Antipolis, France, 1993.

[2] R. Frederick, "nv - X11 video conferencing tool," *Unix Manual Page*, Xerox Parc, 1992.

[3] D. D. Clark, S. Shenker and L. Zhang, "Supporting real-time applications in an integrated services packet network: architecture and mechanism," *ACM Computer Comm. Review*, Vol. 22(4), pp. 14-26, October 1992.

[4] M. Liou, "Overview of the px64 Kbit/s video compression standard for multimedia applications," *Comm. ACM,* 34, Vol. 4, pp. 59-63, 1991.

[5] G. Knight, "OSI-based Management of an Ethernet/ISDN Gateway" *Issues in LAN Management II*, pp101-112. Amsterdam: Elsevier, 1991.

[6] G. K. Wallace, "The JPEG still picture compression standard," *Comm. ACM*, Vol. 34 (4), pp. 30-44, 1991.

[7] D. G. Gall, "MPEG; A video compression standard for multimedia applications," *Comm. ACM*, 33, Vol. 5, pp. 85-110, 1990.

[8] M. Handley, P. Kirstein and A. Sasse, "Multimedia integrated conferencing for European Researchers (MICE): Piloting activities and the Conference Management and Multiplexing Centre," *Proc. JENC '93*, RARE, Netherlands, pp. 18-31, 1993.

[9] T. Turletti and C. Huitema, "Packetization of H.261 video streams," *Internet draft draft-ietf-avt-video-packet-00,* March 1993.

[10] H. Schulzrinne, "A transport protocol for audio and video conferences and other multi-participant real-time applications," *Internet Draft,* 1992.

[11] I. Wakeman, "Packetised video - options for interaction between the user, the network and the codec," *Research Note RN/92/33,* Dept. of Computer Science, UCL, 1992.

[12] D. Verma, et al., "Delay jitter control for real-time communications in a packet switching network," *Proc. Tricom '91,* pp. 35-43, 1991.

[13] I. Chen et al., "A fixed point 16 KB//S LD-CELP algorithm," *Proc. IEEE Int. Conf on Acoust., Speech and Sig. Proc.*, pp. 21-24, 1991.

[14] N. Benveuto et al., "The 32 kbits ADPCM coding standard," *AT&T Tech. J.*, 65, 5, pp. 12-21, 1986.

[15] B. S. Atal et al., "A new model of LPC excitation for producing natural sounding speech at low bit rates," *Proc. IEEE Int. Conf on Acoust., Speech and Sig. Proc.*, pp. 614-617, 1982.

[16] P. Vary, et al., "A speech codec for the European mobile radio system," *Proc. IEEE Int. Conf on Acoust., Speech and Sig. Proc.*, pp. 227-230, 1988.

[17] S. Casner and S. Deering, "First IETF Internet Audiocast," *ACM Computer Comm. Rev.*, Vol. 22 (3), July 1992.

[18] M. J. Handley and S. J. Wilbur, "Multimedia conferencing: from prototype to National pilot," *Proc. INET '92, pp. 483-490*, Internet Society, Reston, VA, USA, 1992.

[19] Y. Chang, "Real-time multimedia communication architecture," *IETF REMCON*, 1991.

[20] M. A. Sasse, M. J. Handley and S. C. Chuang, "Support for Collaborative Authoring via Email: the MESSIE environment," *Proc. ECSCW '93*, September 10-13, Milan, Italy, 1993.

[21] S. E. Deering, "Multicast routing in internetworks and extended LANs," *ACM Computer Comm. Rev.*, Vol. 18 (4), pp. 55-64, August 1988.

Author Information

Peter Kirstein received a B.Sc in Maths and EE from Cambridge U, Ph.D in EE from Stanford U., and D.Sc from London U in EE. He worked at Stanford U, the Centre of European Nuclear Research (CERN) in Geneva, and the US General Electric, in Zurich. He is now Professor and Head of the Department of Computer Science at University College London. Professor Kirstein has been leading research projects in computer communications networks, telematic services, security and multimedia for over 20 years. His current research projects include the PARADISE project in piloting directory services, the PASSWORD and other projects in secured document services, the PODASAX project in ODA applications and databases, and the MICE project in multimedia services. Professor Kirstein is a Fellow of the Royal Academy of Engineering, the British Computer Society, the Institute of Physics, and the Institution of Electrical Engineering.

Mark Handley received his BSc in Computer Science with Electrical Engineering from UCL in 1988. As a PhD student at UCL, he studied novel neural network models and their visualisation. Since 1991, he has been a Research Fellow, working on the RACE CAR multimedia conferencing project and on MICE. His current research interests include Multimedia Systems, especially audio and video encoding and compression, Distributed and Heterogeneous Systems, HCI and graphics.

Angela Sasse has been a lecturer in the department of Computer Science at UCL since 1990. Her main areas of teaching and research are Human-Computer Interaction (HCI), Multimedia Systems and Computer-Supported Collaborative Work (CSCW). Before joining UCL, she worked as a Human Factors Specialist for Philips, who had previously sponsored her PhD studies in Computer Science at the University of Birmingham, UK. She holds an M.Sc. in Occupational Psychology from Sheffield University, UK, and the German equivalent of a Psychology B.A. Before MICE, she worked on the RACE CAR project.