

# MULTIMEDIA CONFERENCING FOR SCIENTIFIC RESEARCH

**Prof. P T Kirstein**

Department of Computer Science  
University College London, Gower St, London, N2 OAR, UK

## ABSTRACT

The talk considers the technological developments in the Multimedia European Conferencing Integration (MERC I) project, which contribute to making multimedia, multi-site, conferencing attractive at this time. It outlines the basic facilities, which are provided in the MERC I system; these include video, audio, shared workspaces, multimedia servers, and mechanisms to tie together both the components and the conferees using a wide spread of network technologies. It discusses the inter-relation between the requirements of the applications, the underlying transport requirements, the facilities which must be provided at the network layer, the existing and planned communications infrastructures of participating organisations, and company policies and needs for security.

## 1 INTRODUCTION

Many forms of scientific research require interaction between research groups. Of course much research requires complex computer computations to carry out the research itself, but these are often not enough by themselves. For research to be done most effectively, it is often necessary to review results interactively by a number of collaborating groups. These reviews must be undertaken within the constraints both of the communications facilities of the participating organisations, and other potential constraints such as privacy in the discussions. These considerations put severe constraints on the nature of conferencing facilities, which can be utilised.

There is a move to specialisation of function and of expertise. It is no longer as necessary to have each specialisation represented at each site; high quality interaction can be achieved over networks. We cannot guarantee homogeneous facilities. Some branches of science and engineering need access to very high speed networks and the power of engineering workstations based on applications using UNIX workstations. For most of their work, many scientists and engineers need only the range of applications available on PCs using Microsoft operating systems, with relatively low speeds of network access? The repositories of data may be specialised to the local user; they may have to be accessible, however, for the remote consultation process.

In this environment, security considerations are of considerable importance. We have introduced facilities for encrypted operation of all the vital tools that we have produced - with robust and scaleable key distribution facilities. There will be some discussion in the paper of the trade-offs of introducing secure working at different levels.

While many of the underlying Standards have come from the ITU-T community (e.g. [1], [2]), most of the standardisation at the levels being considered here have come most rapidly and comprehensively from the Internet Engineering Task Force (e.g. [3]). We will mention in the paper where such standardisation is taking place and its impact on the community.



As a reaction to these various needs of industry, in the series of projects called MICE/MERCI/MECCANO [4] - [6], we have provided a range of facilities designed to address some of these industrial needs. This paper considers the results and accomplishments of that project in the context of collaborative research. We started from the premise that we must provide our facilities in a very heterogeneous environment - both as regards computing equipment and network. We did not address all the needs of collaboration technology; indeed, our primary targets have been more Education and Training. We state both the requirements we sought to address, and how they were addressed.

In the MERCI project, we have started off from the assumption that it is very important to use Standards wherever possible - though these Standards come from different sources. The participating entities may use different communications technologies at the lower levels, though most engineering organisations have now adopted Internet or Intranet communications for many of their activities. We have chosen to base our Audio, Video, Shared Workspace, Multimedia Server, Gateway and Security considerations on the use of multicast procedures over an Internet technology base. For this reason, we have built robust procedures into the all the basic tools to ensure that communications errors in the low-level communications do not destroy performance of the vital activities. We have also progressed advances in deployment of "Quality of Service" techniques, to improve the error rates encountered in the more sensitive components. We have chosen to provide components so that a combination of normal Internet, ISDN, ATM and Direct Broadcast Satellite can all be used as transmission facilities. Finally, we realise that it is improbable that all components can be provided by one organisation. For this reason we have provided powerful mechanisms for the tools to work together on the user's machine via a "Conference Bus"; this should allow good mechanisms for "mix and match".

In Section 2 we indicate the requirements, and in Section 3 consider ~~which~~ we have addressed and with what success. The process of addressing these needs have been a prime example of collaboration technology. In Section 4 we describe how we have used the discipline ourselves. Finally some conclusions are presented in Section 5.

## 2 THE REQUIREMENTS

For maximum collaboration between geographically distributed groups, it is desirable to have similar facilities available between groups as there are inside the groups. This paper is too short to address all the tools that can be provided for good collaboration inside single groups. We must assume that these have been provided, and consider here only the additional facilities required between groups. It is vital to provide access to network multimedia facilities for consultation, conferencing and database access. Groups want a number of specific services; these are indicated below:

- a) **Management and Trouble Shooting:** Users would like to run their conferences and seminars with minimum need to interact with other sites to find out and resolve transmission problems. The set-up and conduct of the sessions should be straightforward and simple. Technical support in one site should be able to support multiple sites.
- b) **Light Weight Conferencing:** Simple digital systems - sometimes with isolated workstations, sometimes in co-ordination with a video projector - using the available network technology.

*A key aspect of this type of activity is the change of roles & what it can be done here we try to do it ourselves.*

*on all the way*



- c) **Range of Workstation Support:** Both PCs and UNIX workstations are required in different environments; it may be very disrupting if users are unable to employ their normal workstations for consultation. It is important also that there be a range of equipment allowing users to upgrade cheaper workstations without going for radical replacement.
- d) **Range of Network Support:** Users often employ the same equipment for local or wide area applications - requiring its de-coupling from particular applications or attached network.
- e) **Video/audio:** Different qualities of service (QoS) are required, depending on the nature of the application, the availability and cost of bandwidth; the audio is the more critical for User acceptance in the consultation process.
- f) **Conference Room Support:** Users may accept a quality of multimedia in workstations lower than for lecture room or conference room presentations; for the latter some practitioners would also like electronic whiteboards - but not too often for concurrent engineering if the other tools mentioned here are available.
- g) **Shared Workspace:** The nature of these requirements depend on the applications area - but it is often a vital feature;
- h) **Introduction of Multimedia Servers:** Some users must record conferences; other users would like to introduce multimedia information, or may be concerned mainly in interacting with the server.
- i) **Audio/video Quality:** Commercial users want multimedia conferencing for general purpose communications in business affairs, but many would accept somewhat lesser quality - if they gain in reachability, speed (of reaction and configuration) and economy. For some engineering applications very high quality is essential; for many it is of less concern. Ease of use, user configurability and interoperability are vital. Mobile access can be important in some fields.
- j) **Security:** In many commercial applications, confidentiality and authentication are indispensable. One needs to protect against unauthorised intrusion, permit transfer of confidential information, allow for verification and non-repudiation of participants; the system must allow only authorised setting of components or provision of session information. More will be said about this area; it is one of the most serious bars to widespread deployment of the tools mentioned in this paper;

**Multiple Locations:** It is often important to have the sessions be multi-way. Sometimes a few locations could participate closely in the same sessions; at others many users would like rather weaker interaction. The use may even be mainly one-way, e.g. for providing information or server access, with minimal reverse communication.

### 3 THE TOOLS

#### 3.1 Introduction

In the previous section we have listed some of the facilities which would ease collaboration between groups. In this section we will describe some of the tools which are available to meet these requirements. Many of these tools have been either developed under the MERCI project, or been evaluated during that project. An important aspect of the tools used is that most are capable of multicast operation. The implications on the underlying network technology is discussed in Section 4.



### 3.2 Audio and Video

There are now many implementations of digital voice on workstations and PCs. Those deriving from ITU-T auspices have often been tied to complete protocol stacks - though the audio and video coding aspects are quite independent. In the Mbone, it has been necessary to tie the codings only down to a packet technology, and have them standardised in isolation. Similar coding standards are used in the ITU-T and Mbone worlds; the nature of errors in packet and circuit networks has led to different implementation choices of the standards. For example, very popular video-coding standards for video conferencing are H.261 [19] and H.263 [20]; these do sophisticated video coding inside a frame, and send the whole frame at frequent intervals. In between sending whole frames, H.261 and H.263 may send the differences between one frame and the next; these differences are called motion vectors. For the comparatively high quality transmission channels assumed over ATM networks, the ISDN or even PSTN, the sending of motion vectors improves quality for a given bandwidth. In the Internet, where packet loss due to congestion in gateways is the norm, the transmission of motion vectors would often increase the bandwidth used without improving quality. As a result, for a given bandwidth the video and audio shown on an ISDN ITU-T workstation looks superior to a Mbone one. When the same transmission is attempted across a concatenation of Internet WANs and LANs, the image may break up and the voice become unintelligible at a much lower error rate for the ITU-T choice of encoding than the Mbone choices.

The pioneering implementations in this field have been the audio tool VAT [21] and the video tool VIC [22] from the Lawrence Berkeley Laboratories; later tools have concentrated on improving particular features of these tools. For higher quality, this includes the use of coding schemes more targeted to higher performance like MPEG-I and MPEG-2; these are still being integrated into the latest versions of the video tools (It will be done under MECCANO). In the case of the RAT audio tool [23] on the MERCI project audio, achieving higher quality has involved adding multiple sampling frequency support, including stereo channels, support for codecs with other than 20 ms unit size, and wide band speech codecs with 16 KHz sampling.

To achieve good quality in the presence of loss, we have taken some of the more popular audio codings, and incorporated a second lower bandwidth coding of a previous packet in a subsequent packet. This has made speech quite intelligible at 50% packet loss rates; such losses would make speech impractical using conventional techniques. Similarly in the RENDEZ-VOUS video tool [24] on the MERCI project, INRIA has provided layered video coding. The different layers are transmitted over different IP streams, thus allowing simultaneous transmission of different qualities of video. At gateways to a lower performance section of the Internet, it is possible to carry only the lower bandwidth streams. Thus it is possible to continue to provide video at lower quality to disadvantaged customers, without penalising others in the same conference with access to higher bandwidth facilities.

Additional components, which are needed to make use of these facilities, are gateways that can transcode and mix audio streams and filter video ones. Such gateways have also been developed in the MERCI project, and are just starting to be deployed. Both in the MERCI project and elsewhere, these audio and video tools have been ported to run under many platforms - both UNIX and Microsoft; often the UNIX variants have not yet been optimised to run under the latest Microsoft operating systems, but they do run in that environment.

So far we have talked of 'workstations'; when these are equipped with video projectors, echo cancelling facilities, multiple camera, multiple microphones, and the ability to switch between cameras, then the same basic data transmission can be employed with groups of



users in a conference room. The extension of the technology to the conference is straightforward; it is quite independent of whether packet switched or circuit switched facilities are used in the transmission. Users may accept a quality of multimedia in workstations lower than for lecture room or conference room presentations; for the latter some practitioners would also like electronic whiteboards - but not too often for concurrent engineering if the other tools mentioned here are available.

From the above, we see that we have largely tackled the problems of audio, video and Conference room support expressed as requirements (e) and (f) in the previous Section.

### 3.3 Shared Workspace

The provision of shared workspace facilities is vital to Concurrent Engineering; the nature of the facilities required is critically dependent on the application. In the Net Meeting [25] suite provided under Microsoft Windows, application sharing is provided. This solution works well for small numbers of participants using ITU-T networks; to extend it to multi-party working over the Internet requires the provision of an underlying reliable multicast facility. It is still a research issue how this should be provided in general; it is not available in Net Meeting. From the arguments of Section 3.2, there are environments where such a system is adequate; it cannot be used in general for the Mbone that is the subject of most of this paper.

A number of tools have been written which provide subsets of the facilities required. The best known WB [26] from the Lawrence Berkley Laboratory allows shared drawing and presentation of pre-loaded material over the Mbone. WB runs only on UNIX systems; a variant WBD [27] from Loughborough U, allows WB facilities to run on PCs using the Microsoft operating systems. Another tool NTE [28], from UCL, allows shared editing of documents; it runs under both UNIX and Microsoft operating systems. Both these have their own versions of reliable multicast at the applications level, and work well in the presence of errors. Several other tools address different aspects of the shared workspace problem; I neither wish to state that the problem has been solved, nor that the two cited are the best currently available. For example Mdesk from the MATES project [29] has similar functionality. They are cited only to show that there are solutions available, which are adequate for many of the problems of concurrent engineering. Thus we have been addressing the Shared Workspace item (g) in the list of requirements.

### 3.4 Light Weight Sessions

For a lightweight session, no explicit session control or membership control is required, but a rendezvous mechanism is required to communicate information about the existence of a session to its potential participants. To this end, the Session Description Protocol (SDP) [13] and Session Announcement Protocol (SAP) [14] were devised by UCL and LBL, originally by the MICE project, and have continued to evolve and be revised by the MERCI project. These protocols are implemented in the UCL Session Directory tool SDR [16]. SDR is now used ubiquitously by the Mbone community. SDP and SAP are currently Internet-drafts, and will be progressed to IETF "Proposed Standard" status during 1998.

In addition to announcing public sessions as a *table of contents*, an alternative mechanism to initiate multimedia sessions follows a telephone style analogy - a user wishes to call another user to start a session immediately. Using the Session Invitation Protocol (SIP) [17], a user may call another user who has SDR running, and invite them to join an existing advertised session or a new private "on-demand" session. The current



implementation of SIP in SDR is minimal, but performs well. SIP, however, allows for relays to provide user-location, firewall transition, etc.

### 3.5 Servers

The provision of a multimedia server is considered vital to the MERCI project. This should have the capability to record sessions, and to introduce material previously recorded. The multicast technology permits the servers to become part of a conference. They will then receive all the data passed in the conference, and may record it if desired. They may also introduce their data into conferences just like a normal conference participant.

Early versions of such servers were created at UCL during the MICE project, and also at many other sites including UC Berkeley [30]. The UCL tool, called MMCR, is discussed in detail elsewhere [31]; it currently has a basic, but robust, functionality; many extensions are still required, however. It can participate in a conference - but it must be started manually; eventually this initiation will be via SIP. All operations can be operated remotely through the Internet. It includes an index track and a timing track for each stream. Another such tool comes from Lulea U [32]; this time the initiation and control comes via the WWW. Moreover Standards for remote operations are still being developed in the IETF [33].

The architecture of MMCR allows for annotation and indexing of each media stream; it allows also for synchronisation between streams. This permits adding further facilities for playout of the MMCR- and its variants in other sites, including fast forward and rewind. For audio and video, this facility is implemented fairly easily by using the timing information on the timing track, and starting at the subsequent I-frame in H.261 or MPEG video stream. For remote presentations, similar considerations apply; there is no state information in the system. For shared workspace streams, the situation is more complex. Each operation by an individual has a specific impact, so that reliable multicast is normally used. On fast forward, the sequence of events must be processed to achieve a consistent state. On rewind, it is difficult to avoid going back to the beginning, and recreating the state information to the time in question. In addition, we envisage an index of the stored material, which can be searched in a sophisticated manner, and tools for editing stored material. All these extensions are still for the future.

The considerations for secure operation of the Server are complex. There is no problem in storing encrypted media streams, though it would be necessary also to store the encryption keys used. We need to consider carefully, however, mechanisms of access control and key management for accessing the data stored. This covers the requirement (h) for servers in the previous section.

### 3.6 Ease of Use

The Mbone tools mentioned above are reasonably easy to use, but they are still separate tools. Various studies have looked at the usability of the tools, e.g. [34]. It is clear that an integrated interface including all the tools to be used - and often tailored to the applications area, is desirable. At least two research networks in Europe are planning some larger scale trials; for this reason they have contracted for more integrated user interfaces, better documentation and easier installation facilities. These are being provided, and will be made more widely available.



### 3.7 Security

The whole question of secure conferencing is complex; a prerequisite is the ability to encrypt the streams produced by the multimedia tools. Eventually this will be done in a tool-independent way using secure IP (IPSEC) [35], the extent to which it will be implemented in IPv4, and the timescale of the implementation, were too uncertain for our purposes. We have chosen to secure the individual tools; this has the advantages that it can be done using the full infrastructure of the RTP transport protocol and of Multicast.

We have arranged for encryption to be provided for the audio tool RAT, the video tool VIC, the shared text editor NTE and the shared whiteboard (WB). The media tools themselves require the provision of information specifying the encryption algorithms used other parameters to describe the encryption and the encryption key. In the version implemented so far for the MERCI tools, the encryption algorithm used is the symmetric single DES [36], and the only parameter needed is the encryption key. In each case, we are using an initiation vector and an encryption key. The developers of the tools have agreed to pass information about the symmetric key in the form of a Pass Phrase; this is conformant to [37], which has been agreed in the IETF. This Pass Phrase is passed through an MD5 [38] module, to ensure true randomness of the encryption key; the first 56 bits after passing through the module are used as the Encryption Key; this mechanism has been built into the above existing media tools. Because the Pass Phrase is 7bit printable ASCII, as are all the other parameters provided by the tools, it is possible to provide the whole Directory Payload, as a printable block.

One aspect of security is the interworking of the tools. This has been demonstrated between the secured VAT from LBL and the secured RAT from UCL, and between the secured VIC from LBL and the secured VIC from UCL. Currently WB and NTE each have only single implementations in common usage; we have not yet investigated security for the recently introduced version of WBD

Another aspect of security is the distribution of the Pass Phrase mentioned above. Mechanisms have been devised using Public Key systems for secured session announcements and invitations. These methods rely on issuing a Group Secret Key (GSK) to authorised participants in a set of sessions. The details of future sessions may then be announced broadly in encrypted announcements - using symmetric encryption for efficiency. The Session Encryption Key (SEK) is then transmitted encrypted with the Group Public Key; as a result, only authorised parties may participate in the conference. The distribution of the GSK may itself be done in different ways; one is the use of secure e-mail. Again all the mechanisms above are being standardised in the IETF.[] 42

Thus we have now covered all the security requirements (j) of the previous section.

## 4 THE UNDERLYING NETWORKS

### 4.1 The Network and Multi-User Working

Fundamental to the tools we feel appropriate for this applications area is the use of IP Multicast and the Internet as a communications medium [7]. The use of the Internet might have required justification five years ago; it seems to require little to day. The ubiquity of the technology inside engineering departments, and the ease with which it can be extended from a local area Intranet to first a wide-area Intranet and then a wide-area Internet, makes it natural choice. Provided that a number of very real network level problems can be resolved, this extension is straightforward. The use of the multicast tools are a natural extension; these make it remarkably easy to extend use to an arbitrary number of



participants; this is just the behaviour needed as it is desired to bring in others to solve outstanding problems. The reality is that the bandwidths of communications available in different parts of a single company vary greatly; the quality of communications resulting is also very variable and heterogeneous. For a specific activity like video conferencing, an ITU-T approach is tenable; a completely different technology like the telephone or of ISDN can be deployed - using MCUs [8] for multi-party communications. For concurrent engineering, - in which the tools used in the engineering process may need to be demonstrated *in situ*, the ITU-T approach does not scale well. It is not even claimed to so scale; the H.120 work of conferencing over the ISDN [9] has been broadened to include use over LANs with H.323 [10]. The way that this could extend for multi-party operation on truly heterogeneous wide-area networks has not really been tackled.

The above should not be construed as meaning that the problems of multicast conferencing have all been resolved; they have not. However, a technology is in place that allows straightforward extension of LAN operation to wide area one. The necessary multicast routers have progressed from research vehicle to inclusion in commercial products. The result is a multicast IP network called the Mbone [11]. The present Mbone as a single entity has very serious performance problems; its universal deployment as a high-performance entity needs large investment and resolution of a number of problems that are still in the research phase. The engineering of Mbone facilities over parts of the Internet, with Unicast tunnelling through others is well understood.

#### 4.2 The Underlying Network Technology

A number of network technologies are being used in the MERCI configurations. These include the normal Internet technology, ATM networks, ISDN networks, Direct Broadcasting Satellites and soon mobile networks.

#### 4.3 ATM Networks

For the normal Internet, the Mbone is tolerably well managed, and permits a total traffic of some 250-500 Kbps to be passed through the Internet. This is not very much considering the extent of multicast usage. For this reason considerable work has been done on limiting the reach of such conferences by *administrative scoping*, and limiting the traffic that is allowed into specific networks by *source route filtering*. One of the techniques that we have applied is to put a large part of our traffic over Unicast VCs as part of the JAMES ATM network. By adding also the administrative scoping and route filtering, we can ensure that only our desired traffic can use these paths. This provides the ability to provide much higher quality conferencing.

The topology used for the MERCI activities is shown below.

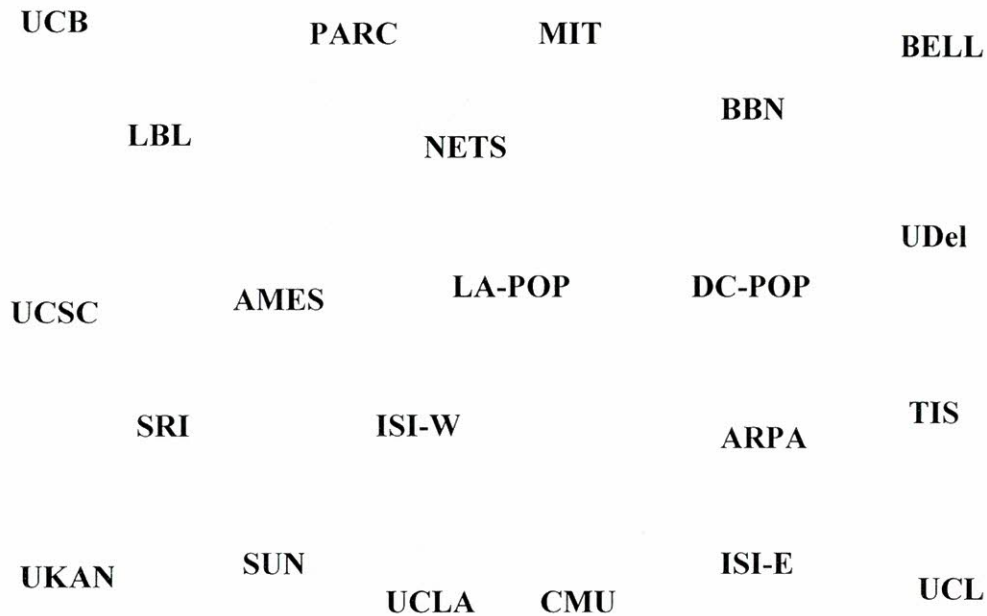


Figure 1. Topology of the JAMES links in the MERCI Network



Besides being part of the MERCI project, and carrying out conferencing in that environment, we are also part of the US DARPA-sponsored CAIRN project. This is trying to bring multicast technology into a really deployable state.

It is not practical to go into detail on the CAIRN project here, but its topology is indicated in Figure 2. Here most of the sites are US based, but UCL is shown in the bottom right-hand corner.



**Figure 2. Schematic of the CAIRN Network.**

In Fig. 2 the thick lines donate 155 Mbps ATM links, and the thin ones 2 Mbps ones

The tools for providing Quality of Service exist; no separate network needs to be installed for their deployment in well-understood environments. Certainly the use of this basic technology allows the deployment with the range of network support and the broad set of diagnostic. Thus for good quality conferencing we are experimenting with various forms of Resource Reservation like RSVP [12].

**4.4 ISDN Facilities**

Although ATM networks give us the highest speeds of access, we are also interested in bringing in sites with much lower speeds. For this environment multicast routing is not yet ubiquitous, so that many sites that wish to participate in Mbone sessions may not be able to do so. This applies in particular to workstations situated at the end of low speed line such as N-ISDN. Such workstations often do not wish to run full Multicast routing, or cannot do so if they are running any Microsoft operating system without a local Multicast router, thus an alternative solution is required.



Another issue is concerned with the heterogeneity, which is intrinsic to the Internet technology. Due to the mismatch of link speeds, the bandwidths available between different sender-receivers are different. In current ITU-T Wide Area applications, each sender transmits data at the same rate to all receivers. With heterogeneous receivers, the source would have to run at a rate that matches the most constrained receiver. For example when N-ISDN is used, the available bandwidth is limited to 64 Kbps (1 B-channel) or 128 Kbps (2 B-channel). Currently most Mbone transmissions use a target rate of 128 Kbps for video and 64 Kbps for audio, precluding simple bridging of Mbone sessions across single 128 Kbps ISDN lines. Thus some kind of rate control - reducing the data rate by transcoding or reducing frame-rate (for video) - becomes necessary.

One possible solution to addressing the above issues is based on the idea of UMbone (Unicast Mbone) which was proposed by Crowcroft et al.[1] at UCL. It suggested a set of components that can be put together to form a relay between the multicast world and a single, multicast-deprived, end-system. Following the UMbone idea, we developed a set of tools which enable users accessing a Mbone session (currently audio and video sessions) from hosts that are not multicast enabled.

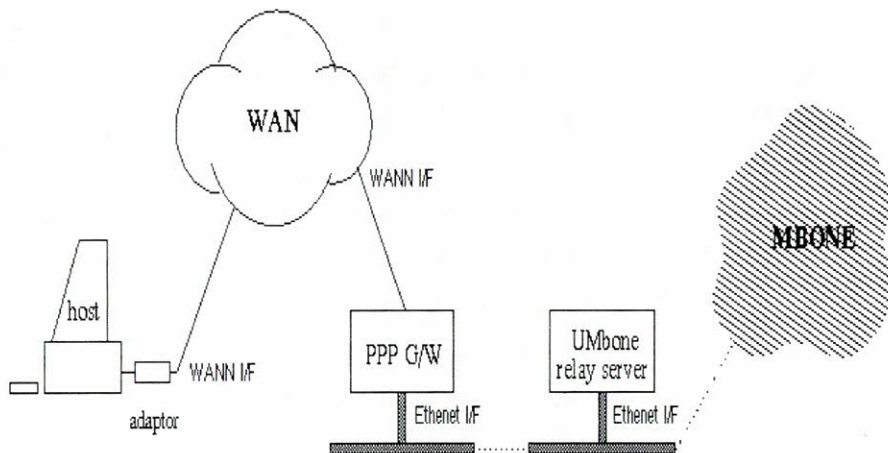
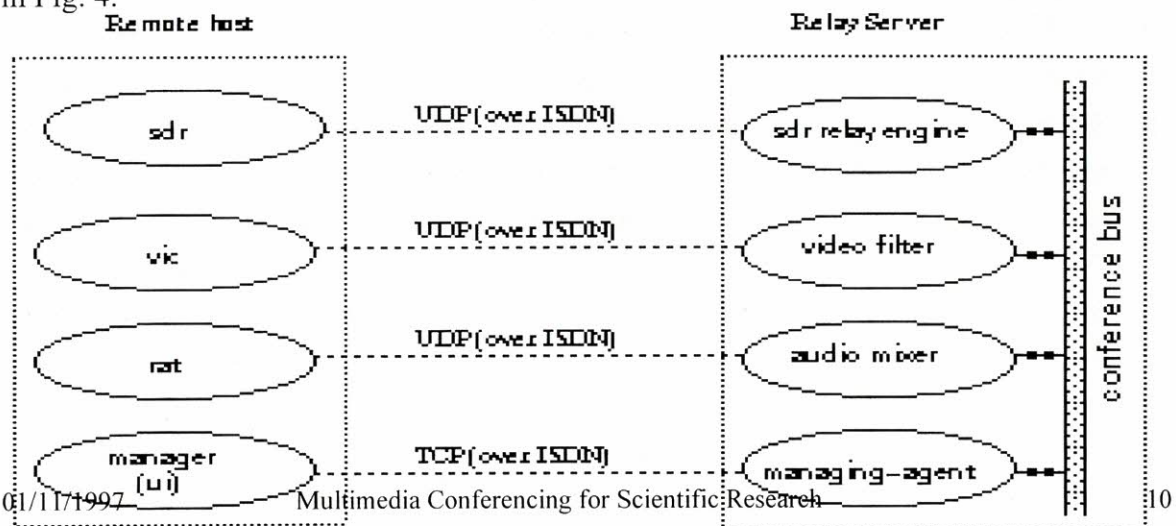


Figure 3. Schematic of Mbone Relays

Full details of this relay are given in [1]. However a schematic of its functionality are shown in Fig. 4.





#### Figure 4 Schematic of the Mbone-Ubone Relay Functionality

Here the relay carries out audio multiplexing, video switching, shared workspace pass-through, video transcoding and audio transcoding. There are a number of complications in carrying out these functions on a multi-user basis, but this is a very useful tool. It is being developed further to deal with the situation where the Mbone configuration is not an ISDN but a mobile gateway.

### 4.5 Light Weight Sessions

For a lightweight session, no explicit session control or membership control is required, but a rendezvous mechanism is required to communicate information about the existence of a session to its potential participants. To this end, the Session Description Protocol (SDP) [13] and Session Announcement Protocol (SAP) [15] were devised by UCL and LBL, originally by the MICE project, and have continued to evolve and be revised by the MERCI project. These protocols are implemented in the UCL Session Directory tool SDR [16]. SDR is now used ubiquitously by the Mbone community. SDP and SAP are currently Internet-drafts, and will be progressed to IETF "Proposed Standard" status during 1998.

In addition to announcing public sessions as a *table of contents*, an alternative mechanism to initiate multimedia sessions follows a telephone style analogy - a user wishes to call another user to start a session immediately. Using the Session Invitation Protocol (SIP) [17], a user may call another user who has SDR running, and invite them to join an existing advertised session or a new private "on-demand" session. The current implementation of SIP in SDR is minimal, but performs well. SIP, however, allows for relays to provide user-location, firewall transition, etc.

Another requirement is for tightly coupled conferences. This is the style expected in the ITU-T series of recommendations. TELES, one of the MERCI partners (represented by TU Berlin and TU Bremen) has been developing the Simple Conference Control Protocol (SCCP) [18] in the IETF, which will provide an IP multicast-based mechanism upon which T.124 semantics can be built. Work on SCCP is less advanced than that on SDP/SAP/SIP, as it started later; running code is not yet publicly available, but an SCCP implementation should be available during 1998.

While it is an over-simplification to say that the use of multicast IP technology has resolved the problems implied in (a) - (d) above, they provide a basis for providing solutions.

### 4.6 ITU Gateways

The alternative to the Mbone tools is the set of ITU-T recommendations - e.g. [1], [2]. Another aspect of the MERCI project is the implementation of a gateway between these two worlds. At the moment the gateway is only a Proof of Concept. A schematic of it is shown in Fig. 5.



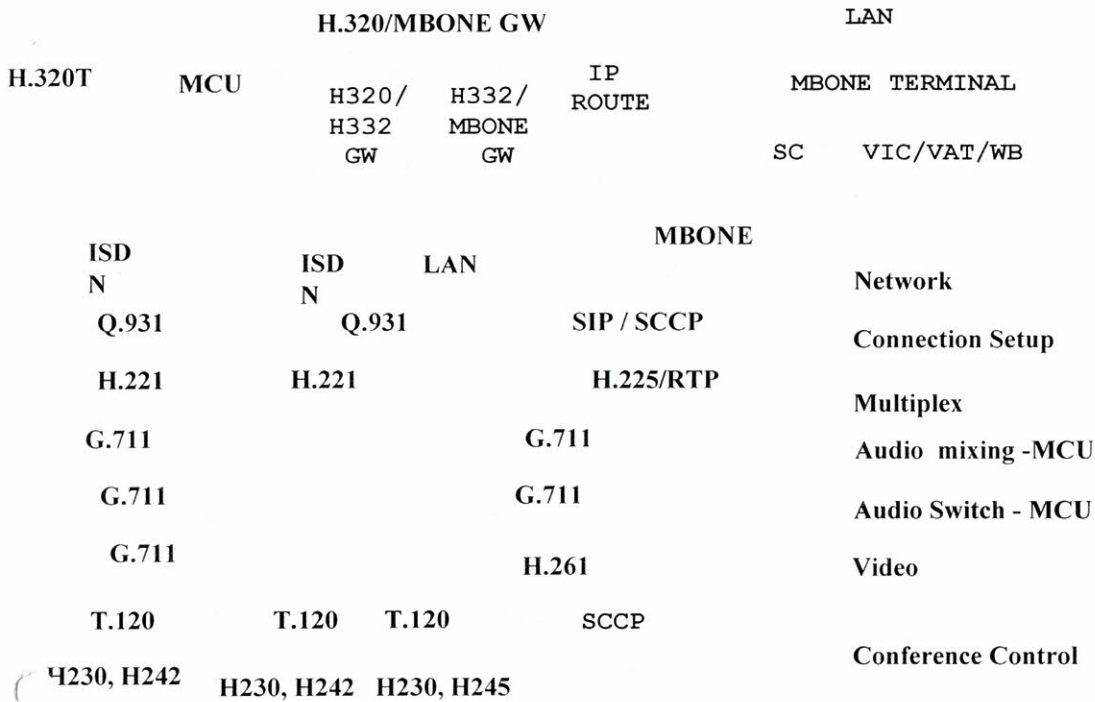


Fig. 5. Schematic of the ITU-Mbone Gateway

#### 4.7 Use of Direct Broadcast Satellites

It is often important to be able to use the Mbone tools in an asymmetric manner - for instance in making presentations. Moreover there are many sites where it is difficult or expensive to obtain high symmetric bandwidth into a site. For this reason we have been exploring also the use of DBS satellites. We have been using a EUTELSAT satellite, with an 2.5m upstation link and 80 cm receiving stations. With these dishes, it is easy to obtain 2 Mbps on the down stream path. The return path is by ISDN or conventional Internet. This configuration is ideal for presentations, where the higher speed allows full motion, high quality ISDN, while the return allows good audio and shared workspace feedback. In our implementations, it is necessary to use DVB as the transport protocol, since this is Standard for Eutelsat. We make the upstation as a member of the multicast group, and have modified the routing protocols so that the system will continue to operate without the correct routing acknowledgements from the receiving stations.

### 5 THE USE OF THE TOOLS INSIDE THE MERCI PROJECT

The organisation of the MERCI project is by itself an exercise in Concurrent Engineering. Its environment meets very well the constraints implied in the bulk of the paper. We have partners with high speed ATM connectivity via the JAMES ATM Pilot; these include partners in Germany (RUS), Norway (Oslo U), Sweden (KTH) and the UK (UCL). Other partners have only normal Internet access from France (INRIA) and Germany (GMD and TELES); From Canada (CRC) access is via a mixture of ATM and Internet. Individuals at GMD and UCL also access their sites using only the ISDN. INRIA has a DBS up-station, and both INRIA and UCL have DBS earth stations. This configuration is sketched in Fig. 1. It fully meets the forms of heterogeneity indicated above. There is a similar mix of

workstations; DEC, others use HP, SGI and Sun workstations are all used by some partners; both Microsoft Windows and NT PCs.

The system is used often for interactive seminars; speakers at one site present seminars to all the other sites. They are seen at all sites, and can be queried at breaks in the lecture. Slide material is used heavily. Reception of the seminars is not always perfect; the problems are more due to poor lighting, faulty microphones, bad echo cancelling and poor lecturing technique than failure with the underlying technology. There can be problems with service quality; extensive monitoring and re-configuration has largely removed technical network problems.

We hold bi-weekly management meetings by video conferencing, in which all participating organisations are expected to attend. Three papers showing the use of the tools are one describing their use for language teaching [39], a large international conference [40] and another the specialised used for surgical demonstrations in the context of a medical workshop [41].

## 6 CONCLUSIONS

The basic framework for widespread use of the tools in Concurrent Engineering now exists. For truly routine reliance on the technology, some further major improvements are needed. These include:

- Integration of the tools into a unified user interface
- Integration of security at conference level rather than tool level
- Provision of reliable and appropriate application sharing tools
- Better management of the Internet for multicast, or good provision of in the relevant section of a corporate network.
- Better documentation and facilities for easier configuration of all the components.
- Better understanding of the sizing of gateways, more automated means of configuring them for specific traffic loads and mechanisms for determining where they should be sited.

Many of these improvements are in hand under the auspices of the MECCANO successor to the MERCI project. Others are being undertaken in the context of commercialisation of the accomplishments in the MATES, MERCI and other projects. The technology has reached the state of being ready for routine use by knowledgeable commercial users. Feedback from such users is required to ensure that the improvements in hand proceed the right way.

## ACKNOWLEDGEMENTS

The work described here is due to many of the people who have been working on the MERCI project. While the list below is not complete, special acknowledgement must be made to the following: John Bolot (INRIA), Peter Feil and Andreas Roszek (RUS), Elfriede Hinsch (GMD), Stefan Braun (TELES), Jorg Ott (TUB), John Andrews, Mark Handley, Vicki Hardman, Isidor Kouvelas, Lambros Lambrinos, Colin Perkins, Angela Sasse and Jin Yang (UCL). It is their work, which is really summarised here.

We acknowledge the support given by the European Commission under the Telematics Project RE1007. In addition, Hewlett Packard, one of the sponsoring partners, has made available a substantial amount of equipment, and the Open View network management facilities to the academic partners. The paper also describes contribution made to the MERCI project from other research projects too numerous to name.



**REFERENCES**

1. Recommendation T.120 (1996) "Data protocols for multimedia conferencing"  
URL: [http://www.itu.ch/itudoc/itu-t/rec/t/t120\\_35511.html](http://www.itu.ch/itudoc/itu-t/rec/t/t120_35511.html).
2. Recommendation H.323 (1996) "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service".  
URL: [http://www.itu.ch/itudoc/itu-t/rec/h/h323\\_39297.html](http://www.itu.ch/itudoc/itu-t/rec/h/h323_39297.html)
3. Schulzrinne H. et al (1996) "RTP:a transport protocol for real-time applications", RFC 1889, IETF.
4. MICE National Support Centre (1997) "User Documentation"  
URL: <http://www.cs.ucl.ac.uk/mice/mice-nsc/tools>.
5. Handley M et al (1995) "Multimedia Integrated Conferencing for European Researchers (MICE): Piloting activities and the conference control and multiplexing centre", in Computer Networks and ISDN Systems, 26, 275-290.
6. Kirstein PT et al (1997) "Recent Activities in the MERCI Conferencing Project", in Proc JENC8.  
URL <http://www-mice.cs.ucl.ac.uk/merci/publications.html>
7. Deering SE (1991) "Multicast routing in datagram internetworks and extended LANs", in ACM Trans. Comp. Syst., 5.
8. ITU Recommendation T124 (1995) "Generic conference control for audiovisual terminals and multi-point control units", ITU, Geneva.
9. ITU Recommendation H.320 (1993) "Narrow-band visual telephone systems and terminal equipment", ITU, Geneva.
10. ITU Recommendation H.323 (1997) "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service", ITU, Geneva.
11. Savetz K et al (1996) "MBONE, Multicasting Tomorrow's Internet", IDG Books Worldwide Inc.
12. Zhang L et al (1997) "Resource ReSerVation Protocol (RSVP)", IETF, ietf-rsvp-spec
13. Handley M et al (1997) "SDP: Session description Protocol", Internet Draft, IETF, ietf-mmusic-SDP.
14. Handley M et al "SAP: Session Announcement Protocol", Internet Draft, IETF, ietf-mmusic-SAP.
15. Jacobson V (1993) "Session Directory (SD), README file, Lawrence Berkeley Labs.
16. Handley MJ (1996) "Session Description Rendezvous, README file", University College London (UCL).
17. Handley MJ (1997) "Session Initiation protocol", Internet Draft, IETF, ietf-mmusic-sap
18. Bormann C et al (1997) "Simple Conference Control Protocol", ietf-mmusic-sccp
19. ITU Recommendation H261 (1993) "Video codec for audio-visual conferences at p x 64 Kbps", ITU, Geneva.  
URL: <http://www-mice.cs.ucl.ac.uk/mice/mmdocs/h261/h261.html>
20. Draft ITU-T Recommendation H.263 (1997) "Line transmission of non-telephone signals, video coding for low bitrate communication", ITU, Geneva.  
URL: [http://www.fou.telenor.no/brukere/DVC/h263\\_wht/h263wht.html](http://www.fou.telenor.no/brukere/DVC/h263_wht/h263wht.html)
21. Jacobson V (1993) "Visual Audio Tool (VAT), manual pages", Lawrence Berkeley Labs.
22. McCanne S et al (1995) "*vic*: a flexible framework for packet video", in Proc. ACM Multimedia'95.
23. Hardman V et al (1995) "Reliable Audio for Use over the Internet", in Proc Inet'95.  
URL: <http://www-mice.cs.ucl.ac.uk/mice/rat/pub.html>
24. MERCI (1996) "Rendez Vous video tool", in MERCI project deliverable D1.

- URL: [http://www.inria.fr/rodeo/rendez\\_vous/](http://www.inria.fr/rodeo/rendez_vous/)
25. Microsoft (1997) "Net Meeting".  
URL: <http://www.microsoft.com/netmeeting/>
  26. Jacobson V (1993) "Whiteboard (WB), README file, Lawrence Berkeley Labs.
  27. Highfield JC (1997) "WBD: A WB compatible Whiteboard", Loughborough University.  
URL: <http://bashful.lboro.ac.uk/ROPA/wbd.html>
  28. Handley MJ (1995) "Network Text Editor (NTE)", manual pages, UCL, 1995.  
URL: <ftp://cs.ucl.ac.uk/mice/videoconference/nte/>
  29. Parnas P (1997) "mDesk : a multicast collaboration tool". Lulea U.  
URL: <http://mates.cdt.luth.se/software/mDesk/>
  30. Holfelder W (1995) "Mbone VCR -video recording on the Mbone", in Proc ACM Multimedia '95, ACM, New York, 237, 238, 545, 546.  
URL: <http://pi4.informatik.uni-mannheim.de/pub/mbone/vcr/>
  31. Lambrinos L et al (1997) "The Multimedia Recorder", Submitted to ICC'98  
URL: <http://www.cs.ucl.ac.uk/merci/publications.html>
  32. Parnas, P. et al (1997) "mMOD the multimedia media on demand system", Lulea U.  
URL: <http://mates.cdt.luth.se/software/mMOD/>
  33. Schulzrinne H et al (1997) "Real time streaming protocol (RTSP)", Internet Draft, IETF, ietf-mmusic-rtsp-06.
  34. Watson A. et al (1996) "Assessing the Usability and Effectiveness of a Remote Language Teaching System". In Proc. ED-MEDIA 96.  
URL: <http://www-mice.cs.ucl.ac.uk/relate.html>
  35. Atkinson, R.: Security architecture for the Internet Protocol, Internet Draft, IETF, ietf-ipsec.arch-sec
  36. ANSI X3.92-1981 (1990) "American National Standards Encryption Algorithm"..
  37. Handley M (1995) "The use of plain text keys for multimedia conferences", RN-95-19, University College (UCL).  
URL: <http://www.cs.ucl.ac.uk/research/rns/RN9519.ps>
  38. Rivest R (1992) "The MD5 Message-Digest Algorithm", RFC 1321, MIT.
  39. Hughes J et al (1997) "Small group language tuition using multimedia conferencing over the JANET Mbone Service" in Proc Networkshop 25, pp 79-82.  
URL: <http://www-mice.cs.ucl.ac.uk/merci/publications/nws25-relate.ps>
  40. Perkins C et al (1997) "Real Time Audio and Video Transmission of IEEE GLOBECOM'96 over the Internet", IEEE Communications Magazine, pp 30-33, April.  
URL: <http://www-mice.cs.ucl.ac.uk/merci/publications.html>
  41. Gevros P et al (1997) "Real-time surgery over an IP/ATM network".  
URL: <http://www-mice.cs.ucl.ac.uk/merci/publications.html>
  42. Kirstein, PT et al (1997: "Specifications of Security in SAP using Public Key Algorithms, IETF Internet Draft, ietf-mmusic-sap-sec-03.txt