# Remote Seminars through Multimedia Conferencing:

# Experiences from the MICE project

Martina Angela Sasse <a.sasse@cs.ucl.ac.uk>
Ulf Bilting <bilting@it.kth.se>
Claus-Dieter Schulz <schulz@rus.uni-stuttgart.de>
Thierry Turletti <turletti@sophia.inria.fr>

## Abstract

*The aim of the MICE project is to pilot multimedia interworking between European researchers, using a heterogeneous hardware platform and existing network facilities. Part of the project brief was to assess cost and benefit of providing a regular multimedia conferencing service for research collaboration, and provide recommendations for the deployment and use of such systems in future. In order to gain gain hands-on experience of both running and using such a service, MICE partners started a distributed International Research Seminar Series, in which researchers and students could participate. Speakers and audiences participated from conference rooms and workstations at MICE partner sites, and increasingly at other remote sites in Europe, the US and Australia. Observations, recordings and comments provided data on which the cost-benefit analysis and recommendations for such a service are based.*

## I. Introduction

The aim of the MICE (Multimedia Integrated Conferencing for Europe) project, which started in December 1992, is to enable interworking between European researchers via multimedia conferencing (audio, video and shared workspace) technology. Rather than develop a new system, the project was to integrate existing heterogeneous hardware facilities at partner sites, such as conference rooms and workstations, and hardware and software codecs. The resulting system has been piloted over existing packet-switched research networks, using both uni- and multicast technology, and ISDN. A detailed description of the project aims and rationale is given in [1].

The project had three overlapping phases: definition, trial and evaluation. During the definition phase, a multimedia conferencing reference architecture has been defined, and facilities required in conference rooms, conferencing workstations and the Conference Multiplexing and Management Centre (CMMC) at UCL [2] have been specified.

During the trial phase, the facilities of all three areas have been continuously improved, and the feasibility of multi-way interworking between the partners and sites in the US was demonstrated successfully in during 1993 at :

- the Joint European Networking Conference (JENC'5) in Trondheim, Norway;
- the Internet Engineering Task Force (27th IETF) in Amsterdam, The Netherlands;
- Interop'93 in Paris, France.

During the evaluation phase, however, cost and benefit of providing a *regular* multimedia conferencing service for research collaboration had to be assessed, and recommendations drawn up for the deployment of such services. In order to gain hands-on experience, partners started to use the technology for weekly project meetings in March 1993. Whilst regular use of the technology for real tasks (distributed software development, collaborative authoring and project management) helped us to identify and address some problems (such as setting of audio levels and bottlenecks in the network infrastructure) early on the the project, a pilot service had to involve potential users from outside the project. Since seminars are a major example of collaboration between researchers, we decided to pilot the technology by setting up a distributed International Research Seminar Series on Communications, Multimedia, Distributed Systems and CSCW, and invited researchers and graduate students to participate as speakers, discussants and audience. The first series of seminars were given on a weekly basis between October and December 1993 (see Table 1), and have continued on a fortnightly basis from February 1994.

| Date | Trans-mitting site | Speaker(s) (Organisation) | Title |
|---|---|---|---|
| Oct 4 | UCL | Ian Wakeman/Jon Crowcroft (UCL) | Congestion Control Schemes |
| Oct 11 | UCL | Graham Knight (UCL) | Narrow-band architecture for ISDN |
| Oct 19 | KTH | Steve Deering (Xerox PARC) | Multicast: State of the Art and Research Issues |
| Oct 19 | KTH | G. Maguire, F. Davoli, F.Reichert, S. Grandhi, H. Tenhunen | Mobile personal computing and communication workshop |
| Oct 26 | UCL | Greg Lavender (MCC) | ISODE Research Seminar: OOSI - Objectified Upper layer OSI protocol stack |
| Nov 1 | KTH | Lars Thylen (KTH/Photonics) | Optical networks and transmission in broadband communications |
| Nov 8 | UCL | Jonathan Grudin (Univ. of California, Irvine) | The CSCW Forum |
| Nov 15 | UCL | Van Jacobson (LBL) | Lightweight session |
| Nov 22 | KTH | Yngve Sundblad (KTH) | CoDesk, the Collaborative Desktop |
| Nov 29 | UiO | Geir Pedersen (Univ. of Oslo) | Distance Education - the MUNIN project |
| Dec 6 | KTH | Per Lindgren (KTH) | The DTM Gigabit Network |

**Table 1: First Seminar Series**

## II. Remote Seminars

This section provides a brief description of the preparation and conduct of MICE seminars.

Since the seminars were multicast using the Mbone [3], on which bandwidth is limited, each seminar was announced on the rem-conf@es.net mailing list several days before the broadcast. The first mail announcement raised considerable interest in the multicast community: the MICE team received a number of enquiries from Europe and the US, asking if researchers from outside project could listen in. The MICE team therefore publicised each seminar on the mailing list, giving time and date, multicast address and port numbers, the speaker's name, title and an abstract of the seminar, and an ftp address for the slides (see next paragraph). At the same time, the seminars were announced using the Session Directory tool , *sd* [4]. Prior to the seminar the speaker provided an abstract of the talk and PostScript or ASCII files of the slides to the moderator at the transmitting site. The moderator will read these into *wb* [5], the shared whiteboard tool through which slides are displayed on remote workstation screens. At the beginning of a seminar, a person appointed as moderator at the transmitting site will introduce the speaker to local and remote audiences, and then pass over the microphone for the talk. Audio is transmitted using the Visual Audio Tool, *vat* [6], and a video image of the speaker is transmitted throughout the talk, using the INRIA Videoconferencing system, *ivs* [7]. Some remote attendees (mainly in conference rooms) are also allowed to send video at a very low rate (less than 20 kbps), which provides feedback to the speaker about attendees and their reactions. Due to problems with network congestion, it is sometimes necessary to adjust the rate at which video is sent during the conference. This is done according to feedback received from remote sites via a control whiteboard (see section 3). In the worst case, if audio quality deteriorates too far due to packet loss, video transmission has to be dropped completely. After the speaker's presentation, questions are taken from local attendees and remote participants, using *vat* [6] as audio tool and i*vs* [7] or hardware H.261 video codecs, and the *speakers* [8] floor tool managed by the moderator of the seminar.

## II.A. Multimedia conferencing tools

All multimedia conferencing software tools used to multicast the MICE seminars are available in the public domain.

*sd* [4] can be used to advertise a session which is to be multicast. It also provides users with a convenient way to join a session, simply by selecting the session entry from a list of advertised events. Audio, video and shared workspace tools are started on the address and with the parameters specified by the advertiser. The advertisement is sent with the same scope as the session itself (i.e. multicast sessions limited to a particular area only advertised in that area). Since resource reservation is currently not available on the Internet, it is very important to know if other events are scheduled to be multicast at the same time. Thus, sd is an informal way to reserve bandwidth by asking other MBone users not to use high bandwidth during the conference announced. *sd* also allows to start each tool with the correct parameters, such as address and port numbers, and the correct packet lifetime (ttl).

The shared whiteboard tool, *wb* [5] from Lawrence Berkeley Labs (LBL), is used as a shared drawing surface in which the speaker can enter all the slides needed for the talk, and display them on local and remote workstation screens. Speaker and remote participants can view the slides, point and draw on them (e.g. to highlight a part or to illustrate a query or suggestion). The speaker can select slides, point to or highlight parts of them, and draw on them.

To multicast audio, MICE currently uses *vat* [6], which supports several encoding modes such as PCM (64 kbps), ADPCM (32 kbps), GSM (16 kbps) or LPC (9 kbps). *vat* also includes mechanisms to dynamically adapt itself to delays introduced in the network, by delaying the playout at remote sites packets. Its graphical user interface is fairly intuitive and can be used without much training or practice. Two different modes of interaction are supported: during the presentation, the tool is used in *lecture mode* (i.e. the speaker's microphone is continuously open); during the interactive question-and-answer and discussion part of the seminar, the tool is switched to *push-to-talk* mode (participants push a button to open their microphone and release it when they are finished). *vat* is a well established tool in the multicast community - for instance, more than 500 people receive multicast of networking conferences such as JENC, INET and the IETF via *vat.*

Both the shared whiteboard and the audio tool can be managed either by the speaker or moderator at the transmitting site. If the speaker operates the tools, some instruction and practise (about 30 minutes) are required before the seminar. If the moderator operates the tools, a "script" needs to be agreed with the speaker and clues given, akin to the "next slide please" in traditional conference presentations. During the presentation, remote participants are usually passive, i.e. they do not send audio or type on the whiteboard.

To receive video and send it when no hardware codec was available, we used the INRIA Videoconferencing System, i*vs* [7]. *ivs* is currently the only video software codec system that conforms to international standards. It includes a H.261 [9] software codec which allows it to converse with hardware codecs such as Bitfield or GPT. The CCITT recommendation H.261 is originally intended for fixed data rate ISDN circuits. A packetisation scheme has been designed and specified in an Internet draft report [10]. It defines how H.261 video streams can be carried over the Internet using the RTP protocol [11]. *ivs* allows to encode video from a frame grabber in the workstation such as VideoPix or Parallax in three different format types: CIF (352x288 pels), and quarter or quadruple CIF. The video output rate is adjustable, and typically is set to values between 10 and 100 kbps.

For floor control, or more precisely for video floor control, we use the *speakers* [8] tool developed by the Swedish Institute of Computer Science (SICS). *speakers* is run at each participant site. The main transmitting site runs this tool with the *moderator* status whereas the listeners sites run speakers with *participant* status. Any participant in the conference is able to request to speak by pushing a button. A list of all participants who request the floor is displayed, and the moderator decides who is allowed to speak next. The moderator can then increase that remote participant's video data rate, and have their image displayed on all workstation screens if they are not already visible.

## II.B. Participating from Workstations

The tools described in the previous section have to be installed on a Unix workstation for remote participants. The hardware needs to support audio, and most workstations are delivered with microphones and speakers. In shared office environments, headphones are often preferred to speakers, and in general, we have found that headphones produce better quality than speakers. Shared whiteboard and video decoding are provided by *wb* and *ivs* - neither require any specialist hardware. If a remote participant wants to transmit video, a video card is needed, such as Sun's VideoPix board or a Parallax card. A reasonably powerful workstation is required to decode multiple video streams, or to encode video [2]. This is a particular concern if the speaker gives the seminar from a workstation, rather than from a conference

room with video codec - Van Jacobson delivered his seminar from a workstation.

### II.C. Participating from Conference Rooms

At all the remote seminars there have been participants at remote conference rooms (at UCL, SICS/KTH, University of Oslo and GMD Darmstadt), with audiences ranging from 5 to 25 participants. Such conference rooms allow participants who do not have a suitable workstations to participate. We also found that some participants prefer to congregate in a conference room, rather than sit by themselves at a workstation, since a conference room environment creates more of a seminar atmosphere.

A conference room needs good facilities for audio. Optimum types and placement of microphones needs the attention of an experienced expert. One could argue that professional audio equipment is wasted since the 8-bit/sample audio resolution of the transmission will not be able to reproduce the quality, but it works the other way around: bad quality input audio, in particular mismatched dynamic levels and noise, will cause the audio coding to use even less than the 8-bit resolution. The MICE partners have invested effort in finding adequate audio configurations.

To make video and a whiteboard visible for a conference room audience, either a large monitor or a facility to project the workstation screen is needed. The MICE partners have implemented different facilities in their conference rooms: back-projected screens with video projectors and light pens or digitiser boards, front projected overhead projector LCD displays, TV monitors for local and remote video. Unfortunately, today there simply are no inexpensive solutions for large, good quality, high resolution, high contrast and mouse-able screen projection facilities. We have discovered that even solutions involving expensive hardware still present problems (see Section III).

Cameras for sending video may be of a less expensive type, provided they have a low noise level - otherwise, the noise tend to confuse inter-frame compression schemes. We found that two fixed cameras, one pointing at the speaker from the back of the room, and one showing the audience from the front or at angle works nicely. Either both cameras can be transmitted, or the local operator can easily switch between the two as appropriate.

## III. Experiences

The main purpose of the seminar series was to get hands-on experience in conducting networked seminars with the multimedia technology available to us.

Our review of the recordings, observations and comments of the first eleven seminars showed that seminars of this kind can further collaboration between geographically distributed researchers and research groups: the seminars gathered a respectable audience, and were viewed as a useful exchange of information, and some of them started a rewarding discussion after the presentation. The discussion even continued by electronic mail afterwards. We feel that the current technical quality of the sound and video is currently at a level which is acceptable, and that the next generation of hardware and network technology, some additional functionality and improved user interfaces will make be able to satisfy the requirements of a wide range of users.

In this section, we summarise the observations made by all participants in the seminars series: speakers, remote and local audiences, and operations and support staff.

### III.A. Speaker's View

At the core of all seminars was the 30-45 minute presentation by one or more speakers. Most speakers were academics working in a computing-related subject. Quite a few were US academics visiting either UCL or SICS/KTH, who could reach an audience in many European countries by being multicast. One of the seminar (by Van Jacobson) was multicast from his office at Lawrence Berkeley Labs, and distributed to some European sites by the CMMC at UCL.

New skills are required to give a lecture to a remote audience. Some instruction and practise is required, and a sound check prior to the seminar to adjust audio levels is necessary. Coping with lack of feedback can be most disturbing. Due to insufficient bandwidth, we often had to keep video rates from remote sites very low (typically 0.2 frames a second with low resolution) or even switched off when the network load has caused a large packet loss. The presence of an audience is very important to almost any speaker who is not specially trained for television. A few seminars were conducted without a local audience and proved very difficult for the speakers, who did not know where to look and felt a distracting suspicion that they were talking into a void. Remote comments on the whiteboard and the presence of video images of the audiences proved to be ways of reassuring the speaker when video quality was low.

A speaker will have the slides displayed through *wb*, either on a workstation screen in front, or

projected on a large wall screen behind. Pointing and further drawing on the slides can be done with the mouse and keyboard of the workstation or with a light pen on the projected image, if available. If there is no light pen, the speaker may still point at the projected wall image and get help from a local support person to put markers on the shared whiteboard. The advantage of this somewhat awkward procedure is that the speaker may more freely address the local audience instead of disappearing behind a workstation screen and lose the focus of the audience.

## III.B. Audiences' View

The seminar series has been run with various types of remote audiences at the different sites: individual users at desktop workstations, remote audiences of 5-25 people in conference rooms, and a local seminar audience in the transmitting conference room.

One of the greatest advantages of using multicast is the scaleability. An added listener consumes very little extra network resources. This means that listening to a seminar from your desktop system provides inexpensive, easy access to remote events. Headphones may be required for desktop participants in shared office environments.

For all participants, audio quality is very important, and the biggest problem is packet loss due to network congestion. We observed that with more than 20% loss, it is almost impossible to understand the speaker, particularly if the language used is not the listener's native language. The use of the distributed whiteboard to display slides helps a great deal with the understanding, but the chopped-up sound resulting from the packet loss is also very irritating and all but the most enthusiastic participants tire and lose interest under such conditions. Remote participants reported that video of the speaker provides higher sense of presence at the seminar than does an audio-only conference. But when audio quality deteriorates due to network congestion, stopping video transmission to relieve congestion and improve audio is preferred.

In conference rooms, since the microphone input is not fed to the local speakers, there is no risk for the wailing sound of audio feedback, and the gain of the microphones may be adjusted as needed. There is, however, a great risk that received audio is going out again through the microphones, causing everybody but the site doing it to get echoes. This is either prevented by an echo canceller or using a mechanism in vat where the microphone transmission is muted by incoming audio from the net, or vice versa.

The local conference room audience is more or less participating in a normal seminar, although there are some differences that are significant. The speaker is addressing not only the people in the room but also the remote audiences, which requires some thought about this from the speaker's point of view. The greatest procedural difference concerns when and how to interrupt the speaker for questions. If the transmitted video image does not include the local audience, it may be hard for the remote audience to follow what is happening. From our experience, the best solution to this is to locally show some remote audience and also the image transmitted. This will increase the feeling of presence and cause "normal" seminar behaviour.

We found that participants experiencing the tools using for the first time are often distracted by their activity on the workstation screen (e.g. the volume indicator in *vat*). We therefore have introduced second screens in conference rooms, so that the moderator or support staff can see the audio tool, but it is not shown to the local audience during the talk. This still leaves the distraction due by activity of the moderator or support person if they have to adjust a piece of equipment or type a response to a query or comment from a remote side.

Finally, in local conference rooms there is a conflict of lighting. The whiteboard, used to show slides and for the speaker to draw on, is usually a projected screen image, which with current technologies is never very bright. Careful attention and testing is required to get the lighting required for video cameras to produce a good image of the speaker who is often close to the projected whiteboard, without losing contrast in the projected image.

## III.D. Moderators, Support Staff and Organisers

One can easily understand why professional video recording teams need all the staff they have: There is a lot of things to take care of to produce a good transmission of a seminar, especially when unexpected events happen, which of course is very often in a prototype environment such as MICE. In future, it should be possible for an experienced speaker to run a seminar without a special technician. So far, support has been desperately needed, and for a speaker new to the environment support is always be needed. At remote receiving sites, a support needs to be present to provide feedback to the sending site technician concerning audio/video quality, packet loss etc. We currently working on monitoring and management tools to provide the moderator at the transmitting site with this information.

In addition to support required during a seminar, there is some preparation and testing to be done beforehand. To coordinate the usage of multicast

addresses use for the different tools (*vat, ivs, wb*), we used the LBL session directory tool *sd*. All necessary set-ups (hardware, software, audio/video equipment and network) have to be done at the transmitting site and all remote sites. We usually start set-up an hour before the start of the seminar. Special care has to be taken on audio set-up. For presenting slides during the seminar we use a shared whiteboard rather than sending video of projected slides for two reasons: the load on the network is reduced and the slides are more readable. This implies that all slides have to be put into the shared whiteboard tool beforehand, so that all participants have them available when the seminar starts. It is of great help to have all slides made available on a software server early enough, so people can get them and keep them locally. In case of problems with the shared whiteboard, these copies can be used instead.

We also use an additional instance of wb running to exchange information between the technicians in charge at each site. This ensures that the main *wb* is used for the speaker's slides only. The best solution is to have a second workstation or at least an extra screen, to keep the control information out of sight of the audience.

We are still working on tools to monitor the video/audio quality at the remote sites, this control *wb* is used to feedback quality information back to the sending site. This way we could best adapt to the sometimes changing network quality. We came up with the following control scheme in case of bad audio quality: The seminar starts sending audio and video. If the remote sites have trouble receiving the audio, they report to the sending site. The technician first reduces the data rate of video. If the audio is not getting better, we usually decide to stop sending video. In case there are still problems we tried to switch audio coding schemes (PCM, IDVI, GSM). From time to time we try to send out video again, and leave it running, if we see that the network problem has vanished.

This easily shows that currently, it is difficult to imagine this conferencing technology is a "walk up and use" environment for speakers and participants, without technical support. A technician will probably always be needed at the transmitting site, at least if the speaker is inexperienced with the seminar tools. It is, however, clear than one technician will suffice to produce a reasonable quality of transmission, let it be that a professional video producer would probably not agree. Maybe we can establish the same kind of relaxed communication paradigm with these new media as in the case of electronic mail. The language style and level of formality of an electronic letter is quite more efficient and comfortable than a traditional letter or even a fax message. Whatever the reasons are for this, let us hope that they can be replicated in networked personal multimedia communication.

## IV. Conclusions and Future Work

Our eleven seminars showed that distributed multimedia seminars are currently possible, using currently available hardware and software tools, and, to a certain extent, the current network infrastructure. Most of the tools and hardware we have used are prototypes or first-generation tools, a fact that bears good promise for the future. The seminars been multicast on the Mbone, giving the opportunity for anyone with connectivity to attend and interact with little extra overhead. The greatest single problem we have encountered has been congestion of the Internet causing unacceptable audio quality due to packet loss. Currently, packet loss is bad enough to turn audio quality unusable, and this in turn jeopardises the entire seminar .

Video transmission is no less sensitive to packet loss. Unfortunately, the compression techniques yielding a good compression also degrades reception quality to a higher degree in case of network data loss. However, in a seminar situation, low video quality is considerably less disturbing to an audience than bad audio. Limited bandwidth and the risk for network congestion calls for tools that concentrates bandwidth usage on the critical media. Good audio is the primary requirement, supported with shared whiteboard.

The participants have been either passive or active, gathered in a remote seminar room or sitting alone at their desktop workstation. Interaction between workstation-based participants and conference room audiences has caused very few problems.

To stage a successful seminar, speaker, moderator and technical support people at both transmitting and receiving sites need to be well prepared. Testing of audio levels, camera position and lighting is essential. Speakers need prepare visual material in advance, and need some practise to master the audio and Shared Workspace tools.

The MICE project will continue in 1994 and we will continue the International Research Seminar Series: it has provided us with a real environment for evaluating the tools and technology, and ideas for how to improve them. The work program for the coming year includes:

*Workstation Components*: We expect that less expensive and higher quality video codecs will be available from major workstation manufacturers. We

will, however, still need to develop schemes for controlled and graceful degradation since all participants in a multicast environment cannot be expected to have guaranteed bandwidth and efficient decoders.

*Audio compression schemes:* Schemes currently used have their origins in circuit-switched networks, i.e. environments with guaranteed bandwidth. Since our experience is that audio is by far the most important medium for seminar activities, we need to invest effort in developing schemes that are less sensitive to network data loss.

*Synchronisation*: We will investigate other factors of multimedia communication such as synchronisation, new protocol standards, shared workspace tools and common protocols for them, evaluation and deployment of new hardware.

*User interface*: The current user interfaces to the software tools are diverse, and the amount of screen space they use create problems when run concurrently. This need to be remedied and support for activity-specific selection of communications tool sets developed.

*Conference rooms:* We will seek to improve conference rooms and specify the most cost effective and flexible configuration for various activities. Of special interest are new developments in display technology. Pilot experiments with developments of tools for other disciplines e. g. medicine and distance learning will be performed.

*Conference Control and Management*: Today the management of the multimedia systems is very awkward, since the tools provided are diverse and often not reachable from outside in a reasonable way. To make fault diagnosis and traffic measurement feasible, a number of activities are necessary: migration to protocols supporting management, modification of current multimedia tools to support management, distribution of the management centres and their functionality, resource allocation and booking schemes.

*Support and Shrinkwrapping:* To gain experience in large scale multimedia interactivity we need to invite a larger body of producers and participants. This will require national software distribution and support centres, MICE-NSCs.

*Multimedia Servers*: We have some rudimentary experience in digitally recording the seminars for later retrieval via the network. We will investigate and deploy recording, archiving, indexing and retrieval mechanisms for audio/video as well as specific data, e.g. minute taking.

*Security:* To be able to restrict participation of a conference in an open network environment, the only possibility is access control by encryption. We will investigate methods for this and for key management. The user interface should hide as much as possible of the security details.

*Traffic measurement, Analysis and Congestion Control:* User behaviour and traffic characteristics are not well understood for integrated broadband communication. Network statistics need to be collected under realistic scenarios of every day usage. MICE provides an excellent opportunity to study properties of packet traffic generated by multimedia applications and provide data to validate traffic models and as a base to develop more accurate models. During our seminars, we realised that the behaviour of the European Internet was fragile and very easy to endanger. We need to analyse the congestion mechanisms to avoid unpermittable stress on the net, but also to furnish network providers with measurements, to make it possible to assess what is needed if they want to provide the relevant network service.

*Applications*: Current applications have been research cooperation and distance education, but mainly with researchers and students in computing-related subjects. We will introduce new applications of multimedia communication, chosen from the environments of the MICE partners to function as interesting examples to validate the MICE approach. These will include MICE adaptation of a 3-D cell structure exploration and remote microscopy.

## V. Acknowledgements

The MICE project (ESPRIT 7602) is funded by the Commission of the European Communities. The authors would like to thank all seminar speakers, who agreed to use experimental multimedia conferencing technology for their presentations, and provided such interesting seminar talks. Other members of the MICE team who provided technical effort at the various sites and made the series possible are *Stuart Clayman, Jon Crowcroft, Mark Handley, Atanu Ghosh, Gordon Joly* and *Peter Kirstein* at UCL, *Knut Bahr, Elfriede Hinsch* and *Hans Mayer* at GMD Darmstadt, *Hans Eriksson* at SICS and *Christian Wettergren* at KTH, *Ronny Nilsen* and *Geir Pederson* at UiO. *Bjorn Pehrson* of KTH and Angela Sasse proposed and organised the first

seminar series. Very special thanks are due to *Van Jacobson* (Lawrence Berkely Labs), the creator of sd, vat and wb, for making those tools available to MICE and quickly dealing with problems whenever we encountered them. Van also encouraged MICE to turn the seminars into a public event, which contributed very much to their success and led to further contacts and cooperation with sites in the US.

## VI. References

[1] P. T. Kirstein, M. J. Handley, M. A. Sasse, “Piloting of Multimedia Integrated Communications for European Researchers (MICE)”, Proc. INET ’93.

[2] M. J. Handley, P. T. Kirstein & M. A. Sasse, "Multimedia Integrated Conferencing for European Researchers (MICE): piloting activities and the Conference Management and Multiplexing Centre", Computer Networks and ISDN Systems, 26, 275-290, 1993.

[3] S. Casner, “Frequently Asked Questions (FAQ) on the Multicast Backbone (MBONE)”, available by anonymous ftp from venera.isi.edu in the mbone/faq.txt, May 6th 93.

[4] V. Jacobson, “SD” README file, Lawrence Berkeley Laboratory (LBL), March 30th 93.

[5] V. Jacobson, “WB” README file, Lawrence Berkeley Laboratory (LBL), August 12th 93.

[6] V. Jacobson, “VAT” manual pages, Lawrence Berkeley Laboratory (LBL), February 17th 93.

[7] T. Turletti, “H.261 Software Codec for Videoconferencing Over the Internet”, Research report No 1834, INRIA, January 1993.

[8] A. Hedstrom, “SPEAKERS” manual pages, Swedish Institute of Computer (SICS), November 9th.

[9] “Video codec for audiovisual services at p x 64 kbit/s”, CCITT Recommendation H.261, 1990.

[10] C.Huitema, T. Turletti, “Packetization of H.261 video streams”, INTERNET-DRAFT, December 5, 1993.

[11] H. Schulzrinne, S. Casner, “RTP: A Transport Protocol for Real-Time Applications”, INTERNET-DRAFT, October 20, 1993.

[12] H. Martinsen [Ed.]: MICE Evaluation Report. Deliverable ESPRIT Project 7602 MICE, 1994.

## Author Information

Martina Angela Sasse has been a lecturer in the Department of Computer Science at University College London (UCL), since 1990. She teaches systems analysis and human-computer interaction. She studied psychology in Germany before obtaining an M.Sc. in Occupational Psychology from Sheffield University in 1986. Previous work includes research on mental models, user interface design, requirements capture, and CSCW. Her main research interest is design and usability issues of computer-mediated communication and multimedia systems. She is the project manager of MICE.

Ulf Bilting is a Research Engineer at Dept. of Teleinformatics, Royal Institute of Technology (KTH) Stockholm. He received his education in computer science at Chalmers University of Technology and the University of Gothenburg. He has been working for the national academic network in Sweden and the Nordic countries since 1980. His current research interests are in networked education and multimedia telecommunication applications.

Claus-Dieter Schulz is a researcher on multimedia systems in the Department of Communication Systems at the Computer Centre of the University of Stuttgart (RUS), Germany. He obtained a Ph.D in theotical physics and worked in laser physics, chaos theory, synergetics and neural networks before joining RUS to work on MICE and other multimedia research projects.

Thierry Turletti received a B.Sc. in Computer Science and Digital Signal Processing from University of Nice-Sophia Antipolis in 1990. Since 1991, he has been a Ph.D student at INRIA Sophia Antipolis specializing in multimedia systems and is working on MICE project since 1992. His areas of interest include audio and video compression and congestion control algorithms.