Interacting with Multi-media, Multi-user Systems: Observations on multi-media conferencing tools

Martina Angela Sasse
Department of Computer Science

Andrew C. Fentem
Ergonomics Unit

University College London
Gower Street
London WC1E 6BT
tel: +44-71-380 7212
fax: +44-71-387 1397
email:a.sasse@uk.ac.ucl.cs

ABSTRACT

Interactive multi-media, multi-user systems systems are becoming increasingly common. Such systems have the potential to change the way in which we work and communicate, since they create the basis for applications in teleworking, distance education and global collaborative projects. Guidelines and design principles aimed at producing effective user interfaces to such systems have been put forward in the CSCW and multimedia literature. This paper introduces some of these guidelines, and examines whether they address the usability problems observed with such systems. The observations were obtained through a small exploratory study with a multimedia conferencing system. We identified several of the problems encountered by the users in the study which were not addressed by the design principles and guidelines. In conclusion, we put forward a list of items which should be considered in the design of user interfaces to multi-media, multi-user systems, and incorporated in human factors guidelines.

1 Introduction

With the rapid spread of global communications networks such as the Internet and ISDN, the use of interactive multi-media, multi-user systems is becoming increasingly common. Rather than being confined to executive video conferencing suites, the technology is moving onto users' desktops, putting access to global communications at individuals' fingertips. Such systems have the potential to change the way in which we work and communicate: they create the basis for applications in teleworking, distance education and international collaborative research. Whether they will realise that potential will obviously depend on them being accepted by a critical mass of users. Given the failure of sophisticated CSCW applications in the past (e.g. Francik, 1991; Grudin 1990), that acceptance should not be taken for granted. Rather, it will depend on whether (a) the technology is seen to support organisations effectively in performing their work; and (b) whether individual users find the systems accessible and usable. System vendors and the European Commission, who is keen to encourage the uptake of broadband communications in Europe, are aware of this. Guidelines for user interfaces have been developed under programmes such as ESPRIT and RACE, and design principles have been put forward in the literature. In discussion with the, however, some developers of advanced multimedia systems indicated that they did not feel these guidelines proviided relevant

and spplicable knowledge to help them make design decisions in user interfaces to complex interactive systems.

2 Multi-media conferencing

Multi-media conferencing systems are one example of multi-media, multi-user systems. They provide not only video and audio links between the remote parties, but also support some kind of shared workspace such as shared drawing and editing tools. Generally, modern multi-media conferencing systems are based on the idea of window-sharing, first implemented by Crowley et al. (1990). Such systems can be used for desktop conferencing, usually on a high-end workstation or from specially equipped conference rooms and lecture theatres. The user can work in windows representing private workspaces and shared workspaces simultaneously. For example, the user could edit a private document in a word-processor window, whilst at the same time refering to an image in a shared sketching tool window.

The main criteria for the selection of the specific system that was evaluated was that the system is in relatively common use, thus ensuring the availability of expert users, and ensuring meaningfulness of the results to a wide audience.

The system used as an example is the one piloted by ESPRIT project MICE (Multi-media Integrated Conferencing for Europe). This project aims to provide facilities for interworking between a number existing technological developments in area (conference rooms and desktop conferencing, packet- and circuit switched networks, hardware and software codecs, uni- and multicast (for a detailed description of the technology solution, see Handley et al., 1993). The intended users are European researchers: it is hoped that such facilities will improve collaboration between them and also link them with research establishments in the U.S. The technology has been used for weekly project meetings since March 1993, and for distributed international research seminars since October 1993 (Sasse et al., 1994).

The MICE project uses a number of multi-media communication tools in order to demonstrate applications of the integrated services system. The system consists of the following public-domain software components or tools, each providing a different medium or communication channel:

- **VAT** (Visual Audio Tool) an audio teleconferencing tool that allows users to conduct multi-party audio conferences over the internet (international data communications network). This tool was written at Lawrence Berkeley Laboratory, California, U.S.A..

- **IVS** (Inria Videoconferencing System) - an audio and video conferencing tool. In this particular experimental configuration this tool was used to provide a multi-way video link over the internet. This tool was written at INRIA Sophia Antipolis, France.

- **NV** (Netowork Video) - a video tool written by Ron Frederick at Xerox PARC, California, U.S.A.

- **WB** (Whiteboard) - a distributed whiteboard facility allowing multiple users to write or draw simultaneously onto a shared workspace, or port existing pre-prepared documents into this workspace. This tool was written at Lawrence Berkeley Laboratory, California, U.S.A..

All of these tools operate on the principle of shared-windows based on the X-Windows system. The communication networks used generally depend upon the circumstances of the conference (whether local or international). They are usually packet-switched networks (often using multicast techology) or

ISDN. The configuration of the hardware and software often has to be tailored to the specific requirements and facilities available for a conference.

The tools are all run on Unix workstations with, ideally with high-resolution colour monitors suited to presenting the complex visual information that these tools generate. Audio input and output is available from a free-standing microphone and loudspeaker combination, or using a headset. The input of video information is via a standard video camera interfaced to the workstation using a digitiser board (see Figure 1).

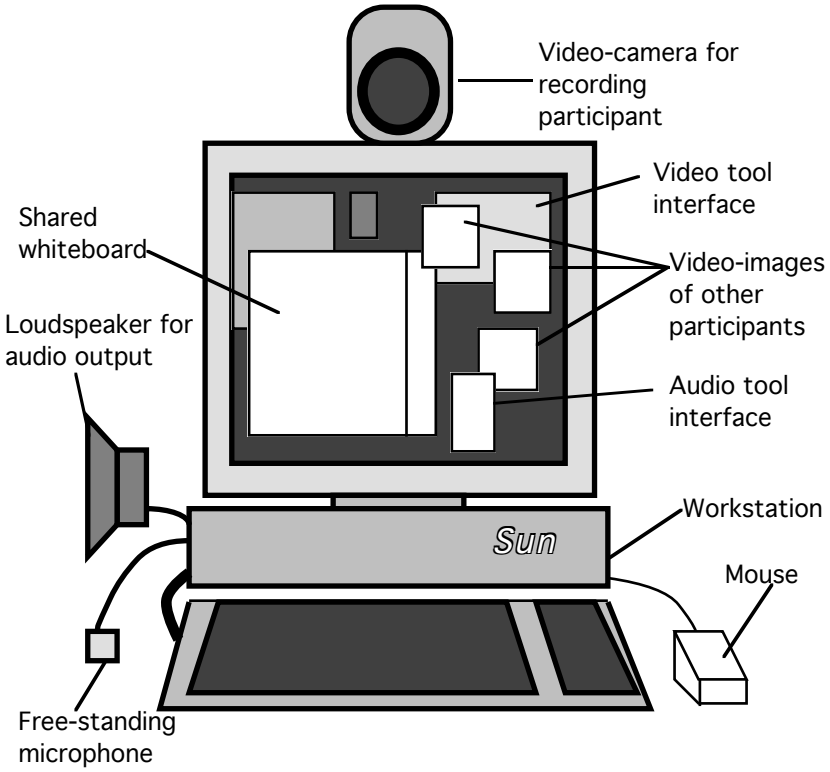

Figure 1 Components of a multi-media conferencing system

3 The Study

3.1 Objectives

The aim of the study was to establish what substantive human factors knowledge is required in the design of the multi-media, multi-user systems. This objective can be approached in two ways: criticism of existing design recommendations and the proposal of new ones. Both of these approaches were taken. Since they are based on an exploratory study with a small number of subjects using one particular system, conclusions drawn from the results and suggested solutions should be regarded as specific and tentative, and further research will be required to assess their validity. Since we feel that some important issues are currently not adressed at all in the design guidelines or the literature, and the technoloy is spreading rapidly, we felt it was both important and legitimate to put them forward at this stage.

The use of multi-media conferencing systems can be viewed from many different perspectives. For the purposes of this study, the position of Kraut and Galegher (1988) was initially assumed, as this approach is more amenable to conventional human factors evaluation methods. This approach allows

the study to concentrate on the ability of the system to facilitate the communication of task-oriented information, thereby aiding performance of a specific task. It is important to acquire knowledge about the nature of the conferencing tasks performed using the system. This allows task analysis and provides an outline of user requirements. Acquisition of user performance data (errors and problems) and user attitudes is important to the evaluation.

The opportunity was also taken to gather data concerning the more social aspects or 'Group Outcomes' (Pinsonneault et al, 1989) of the usage of the system. Finally, we wanted to obtain data which would give us an indidcation which issues were not addressed by current design principles such as WYSIWYS (Elwaert-Keys et al., 1990) and seamlessness (Stefik in Foster and Tatar, 1988), and exisiting guidelines for the usability of integrated broadband communication systems, such as RACE Guidance (1993).

The following approaches were chosen in order to achieve these aims :

a) Study the usage of the multi-media conferencing system. Identify and analyse the conferencing tasks that the system was being used to support.

b) Identify users' (both novices and experts) problems with the system and their attitudes towards its use

c) Identify whether or not the system conformed to any existing relevant ergonomics design guidelines, and assess the validity of these guidelines by attempting to gauge whether these conformities or non-conformities had any impact on the usability of the system.

3.2 Trials Session

The studies were divided into the following sections :

An initial, informal pilot study was carried out consisting of the observation of the system where the conferencing task was the conduct of a small research group meeting, the semi-formal weekly conference of the MICE project.

The second pilot study consisted of the design and distribution of a pilot questionnaire by electronic mail to experienced regular users, in order to discover details of how the current system was being used, and what usability problems were being experienced.

The main formal study consisted of the observation (and video-recording) of a number of user-trials of the system. The novice users involved were asked to perform a realistic and coherent, but varied set of conferencing tasks, with one other remote conferee, in order to measure how well the system would support a real meeting between users with no specialist knowledge.

3.3 Questionnaire and structured interview

Evaluation criteria
The subjects' qualitative assessments were requested in terms of general criteria such as attitude, effectiveness, flexibility, learnability, usability, and functionality. They were also asked for diagnoses of problems they encountered and suggestions for improvements to the system.

Usage questions
Questions regarding usage patterns or in other words, how the system was being employed by the users, can be categorised according to the G.C.S.S. (Group Communication Support Systems) framework suggested by Pinsonneault and Kraemer (1989). Note that only the issues addressed by the framework, and relevant to this study, are listed in Table 1 below.

Contextual variables	**Examples**
Personal factors	Background Abilities Relevant experience Individual motives for using the system.
Situational factors	Reasons for group involvement
Group structure	Group size Interpersonal distance Degree of anonymity within the group.
Technological support	GCSS hardware and software
Task characteristics	Complexity Nature
Group process variables	**Examples**
Communication chars.	Task-oriented communications Non-verbal communications Efficiency of the communication Length of conferences Frequency of the conferences Alternative media available
Interpersonal characteristics:	Co-operation Domination
Task-related outcomes	Quality Cost e.g. time.
Group-related outcomes	Satisfaction i.e. willingness to use the system to communicate with the group in the future.

Table 1 Issues raised by the GCSS impact framework

All of the example issues listed in Table 1 correspond to questions contained within the questionnaire. In theory, the results should specify how the communication process and outcomes are impacted by the specified system and by the specified context in which it is used.

3.4 Other questions

The usability issues identified in RACE Guidance (1993) as being associated with conferencing are **Call control, Call set-up and close-down, Authentication, Addressing, Call Ownership / Call handover, Collaborative manipulation of information, Costs, Service dropping / Alteration and Fallback,** and **Feedback**. A question was designed in order to assess the significance of each issue.

It was thought that the significance of issues such as seamlessness and WYSIWIS would be apparent from the evaluative comments made by the users, and no questions were asked explicitly about such issues.

4. Results

The parameters identified in the design of interfaces to multi-media, multi-user systems presented in this section are based upon the results of the analysis of the experimental observations, structured interviews, and returned questionnaires. This section also incorporates knowledge from HCI and ergonomics literature, and from the authors' own experiences and observations during the two-month trials period.

4.1 Desktop multi-media conferencing system usage

The observations regarding the usage of the conferencing system can be categorised as follows :

4.1.1 Contextual variables

The number of participants involved in desk-top conferencing sessions is generally less than 5, and usually just 2. All of the participants were usually familiar work-colleagues. A very broad range of conferencing tasks were carried out, however the most common were informal talks and listening to lectures. The conferees were often in different countries and time-zones. It was found to be important to have both shared whiteboards for general conferencing tasks such as sketching, and presentation of documents and agendas.

4.1.2 Group process

The length of conferences varies greatly, for the regular users this variance was between 1 minute and three hours. The short length of some of these conferences suggests that conferencing tasks were very different, and would therefore have completely different requirements: For example the very short conferences would need very low 'overheads' in terms of conference set-up time.

Addressing the other participants is generally restricted to calling someone by their name. Familiarity with the other conferees and the context of conversation provides enough cues for participants to be able to tell who was talking, but this would not necessarily be the case in a large meeting between unfamiliar people.

Non-verbal communication is used, in so far as conferees still exhibited gestures etc., but this is interpreted very well, especially by novice users, in view of the frequency of pauses and misunderstandings that they all experienced. This probably explains why most regular users avoid discussing 'tricky' matters over the conferencing system.

4.1.3 Task and group-related outcomes

Most of the trials subjects, and the regular users seemed to be satisfied with the ability of the system to support the conferencing task as few thought that it would take much longer than face-to-face interaction and many thought it was 'fun'. Most of the subjects experienced some affective costs; usually frustration at not being able to convey some information because of the poor video quality.

The social interaction involved in using the system is fairly satisfactory; no subjects or regular users felt unduly uncomfortable or embarrassed. Few subjects felt that they 'knew' the other participant after the conference, although they all felt that it was useful for meeting unfamiliar people. The system is capable of supporting a certain amount of social or relationship-building interaction. Whether this relationship-building is adequate would largely depend upon the social fabric of the organisation or conference group, and how important social bonding is to achievement of the overall organisational goals.

Users' confidence in the system rapidly increased during the course of the trials. The degree of the system's acceptability is illustrated by many of the trials subjects and regular users claiming that they

would use the system to communicate with someone in the next office. However, most said that it would depend on the particular social context. The video-channel does appear to play a significant part in increasing the 'satisfactoriness' of the conference.

4.2 Design recommendations

The recommendations documented in this section are be generalised into the following categories of heuristic principles for general multi-media conferencing system design. Any terms in quotation marks represent concepts coined in whole or part by the authors. These following proposals are being offered as hypothetical solutions that need to be implemented, tested, and validated by future research.

4.2.1 "Independence"

The ability to use many media **simultaneously** e.g. workspace and audio, means that the **input** mechanisms for each of media should be as **independent** from one another as possible. Therefore, as far as is technically possible, all control of the audio channel should be vocally, and of the workspace manually.

This principle extends to cursor control, i.e. the keyboard should control the text cursor, and the mouse the sketching cursor. In essence, a range of methods of input information for the same conferencing task should be provided, and each method should be independent of the availability of other input devices.

This would allow for simultaneous use of the media, and provide a stronger link for the users between the input devices and their effect.

4.2.2 Reception feedback

The participants in a conference should be able to tell how well the output from their media input devices (audio, video etc.) is received by the other remote participants.

For example, the unreliability of the audio link means that the users need some feedback to the user on how they are received at the remote site. The microphone level indicator on the audio tool simply indicates what level the user is speaking at, not how they are received by the other participants. The latter information is much closer to confirming the achievement the users goal task of communication than the former.

4.2.3 Participant and speaker information

In all conferences a list should be available of all participants and brief details e.g. organisation they belong to. To aid participants, especially in conferences consisting of more than one unfamiliar person, it should be salient who is speaking at any one time, as this was occasionally found to be problematic during the use of the system. The names of the video source (as opposed to the internet address) should be positioned close to the appropriate video image . Some reinforcement of the received audio signals and identification of the source should be provided. For example, the perimeter of the appropriate video image and accompanying name could be highlighted.

4.2.4 Seamlessness

During the trials the users found that they were very constrained in the way that they could convey information, so often had to use slow or inappropriate means, for example holding maps upto the video camera.

The principle of seamlessness applied in very constrained and limited ways would allow for conferees to express themselves more fluidly, spontaneously and naturally as follows :

a) "Window seamlessness"
Facilities that would enable any format of document to be very quickly introduced to the conference, on demand. It should be possible to 'grab' other windows, move them into the shared workspace, and modify/annotate them using the whiteboard. For example, images could be 'grabbed' out of the video-window, or from personal workspace windows.

b) "Event-capture seamlessness"
It should be possible to capture conference events in any format that the user requires i.e. full video, printout, copies of windows, postscript files etc..

Other example formats are: 'live recording', 'snap-shots'(single screen frames) and 'snippets' (short excerpts of the conference) of the conference. These formats also have the advantage of supporting long-term memory by providing the information and the **context** (Wickens 1992) in order to trigger the recall of unrepresented information. This recording could be controlled automatically.

4.2.5 Video-window management and content

The multi-media conferencing terminal screen was always very cluttered with windows, especially from the video-tool, making it difficult to find information. To remedy this, there should be management of the windows. The video windows should be easily resizable, perhaps automatically, with the current speaker's window increasing in size, possibly to obscure the other video windows. Maybe only one video-window should be provided at a time as the subjects found one hard enough to concentrate on; the other participants could be iconised when not participating, to save screen space, but serve as a reminder who is present at the meeting.

Users felt that he benefits of the video picture of the other conferees lessen after the first few minutes of a conference. Just as during a normal conversation, users should be able to choose what they can see, and how well they can see it e.g. objects of interest or 'around' other offices to look for other people etc. Issues of privacy (and possibly security) would be raised by such a facility.

4.2.6 Shared workspace functionality and flexibility

The whiteboard, as it stands is a very simple, typical computer aided sketch-pad which formed the focal point for the multi-media conferences. The relative positioning of the communication channels should be arranged so that the shared workspace is at the centre of the terminal screen. This was the subjects' preferred position, with the video image of each of the other participants around the outside of the whiteboard. This arrangement would have the benefit of enabling the user to focus on the most task-oriented part of the system. At the same time the user would be able to use peripheral vision to monitor any "movement" or change of circumstances in the remote sites, as peripheral vision is particularly good at detecting movement (Wickens, 1992).

Use of the whiteboard is very intuitive and has no features that are different to other drawing tools and/or hard to use, and works reasonably well as it stands. However, as it stands it does not support many typical conference tasks due to this lack of functionality.

Functionality that was found to be lacking in the evaluated system includes :

a) Saving and printing of the whiteboard session at any point during the conference is essential. Several possible ways of doing this are listed below :

* Make short real-time recordings of the session. This would have the additional benefit of recording the context of the information recorded, which aids retrieval from long term memory (Wickens, 1992).

* Save the whiteboard to a file.

* Produce a printout of the whiteboard. This has the advantage of being referencable during the rest of the conference.

b) Grouping of objects on the screen should be possible so that they could be deleted or moved.

c) Modification, or at least movement the work of other conferees so points could be ordered (e.g. into a list of actions from the meeting) should be possible, though controlled in some way. At least one person, the minute taker or chairman, should be able to do this.

The workspace should provide as much (or as little) functionality as the user can cope with or desires. Even discussing and modifying a complex 3-D drawing, not all of the conferees will need the tools and expertise to manipulate it, only view it, so should not even be offered the functionality, as this will make their interface more complex. The software for the workspace should be available in a modular form, all modules mutually compatible. This view is supported by Sasse et al (in press).

4.2.7 Conferencing modes : Conferencing task and context oriented design

The facilities required of the conferencing system was found to vary greatly according to the main task of the conference. Hence, the design of multi-media conferencing systems should either be very tightly constrained to support only specific types of conference, with certain types of group. Alternatively, the system could be provided with a selection of different modes of use. The mode of use selected would determine the following parameters of the design, identified during the study:.

- **Metaphor** and associated behavioural options e.g. 'pop-in', 'meeting room', 'organisation building', 'whispering' etc.
- **Screen-style**.
- **Hardware configuration** e.g. all participants have terminals or share one large screen; headsets or loudspeakers for audio output.
- **Verbal protocol** (floor control) e.g. multi-way (open floor for a one-to-one informal conversation), turn-taking (for a formal meeting) etc..
- **Set-up overheads** i.e. time and effort to set-up and amount of screen-space. For example, for a short chat, all of these factors would be low.
- **Bandwidth** i.e. video and audio quality
- **'Synchronicity'** of conference i.e. how much interaction required in the meeting. e.g. observing a lecture passively does not require real-time synchronised audio and video.
- **Media** available e.g. if one media can not be used.
- **Anonymity** e.g. if voting task involved.
- **Iconisation** e.g. other conferees could be represented as names, icons, video-icons, or full video images.
- **Awareness** required of other users' actions and their consequences. For example where they are pointing or their last meaningful_action.

4.2.8 Media integration

All of the media should be set-up with one 'call' and share one interface. This would reduce the number of windows on the terminal screen and hence the apparent complexity of the interface and difficulty for the users to find useful information. Other advantages identified by the study are:

- Take up less screen space
- Cause less confusion for user
- Reduce visual load on the user
- Reduce the system set-up work load
- More consistent user-machine dialogue
- Audio is synchronised with the video.

4.2.9 Real-time and synchronised communication

All regular users and subjects felt that delays in the system caused problems. The delay of the video-image behind the corresponding audio signal had a significant effect on usability. This delay made non-verbal communication very ineffective.

One of the most satisfactory aspects of the whiteboard was that the user input was registered so quickly that it gave the impression of both parties writing simultaneously. Ideally, this principle should be applied to the other media. This will at least benefit novice or irregular users who will not have to become accustomed to a new style of interaction. Video quality was so poor that many subjects questioned it's usefulness. Also it was concluded from discussions with subjects that poor quality, delayed video could 'sabotage' communication by presenting images either inappropriate to the context of conversation, or ambiguous, or distracting.

The conferencing task communication processes could be considered to have a certain **'maintenance factor'**. This represents a level of quality below which the media will not support communication, or even make it degenerate (for social psychological reasons). Some of the media design factors that determine this level are: video resolution and frame rate, audio quality, and media synchronisation (i.e. all up-dated at the same time).

4.2.10 Workspace awareness - how and why things happen

Many of the users experienced problems with identifying the other participant's contributions, especially while trying to perform other concurrent tasks. They complained of having to scan the whiteboard for the most recent additions. This problem should be solved by making the most recent additions (from other participants) more salient to the user. This could be achieved using a host of devices for grabbing visual attention such as colour or flashing.

All conferees being able to see each others' cursors and actions (for example see Ishii and Miyake, 1991) **may** be undesirable if there were many participants, but in a one to one meeting it would lead to more awareness of the other person's intentions, and allow the user to anticipate the position of new input. The intention to change pages on the whiteboard should be indicated or agreed upon first, and WYSIWIS should be strictly enforced here, to avoid the confusion caused in the trials.

Depending upon the nature of the meeting it may be desirable to assign colours to different participants. This would be acceptable to most users as they do not seem to make much use of the availability of different colours, however I think that would be best left to negotiation during the meeting rather than constrained by software as there are also advantages of anonymity of contribution sometimes during a meeting.

In general, awareness of others helps you to anticipate their actions. Meaningful actions in the workspace, such as changing the whiteboard page should be indicated to the other participants before it happens or after agreement. Awareness of less important, or 'meaningful' actions by other conference participants e.g. hand position may not be as useful. The relative benefits of the knowledge of different categories of multi-media conference behaviour by other participants needs to to be clarified by future research.

4.3 Evaluation of extant human factors knowledge

Summarised in this section are the findings of the study with regard to extant human factors knowledge involved in the study.

4.3.1 Theoretical frameworks for ergonomic design

Using the GCSS framework (Pinsonneault et al, 1989) as a basis for the design of questions during the evaluation studies proved to be very effective in documenting all of the relevant parameters of context,process and effectiveness of the conferencing system's use.

4.3.2 RACE Guidance (1993) usability issues for conferencing tasks

The results of this study have shown that the following design principles in RACE Guidance (1993) are valid and applicable to the design of multi-media conferencing systems:

- **Call control**
  The RACE Guidance (1993) principle of feedback, that the "users should be informed of how they are received by other parties", was identified to be one of the more significant deficiencies of the system by the users.

- **Call set-up and close down**
  The conference 'call set-up' procedure was identified during the trials as a very problematic part of the system. This state of affairs would probably have been predicted from the RACE Guidance (1993) principles.

- **Addressing**
  The facilities offered by the system that adhered to the principles in this section of RACE Guidance (1993) were system facilities that were observed to be particularly useful. These facilities were group addressing and the directory. These groups have a pre-established set of permanent members, the name of the group usually associated with a specific task. Each name in the directory can be used as a means of quickly calling up a whole working party, with the combined expertise to tackle a specific problem.

- **Collaborative manipulation of information**
  The principles as outlined in this section of RACE Guidance (1993) are predominantly pertinent to the design of the shared workspace. This section states that the principles of concern here are flexibility, controllability, and salience.

  The justification given for the inclusion of 'salience' in this list is that "in co-operative workgroups users require knowledge of both their own and other users actions and their consequences". This requirement which could be termed 'awareness' was neither conclusively rejected or accepted by the results of the studies. The reason for this is that this approach should be constrained in so far as users only need to know **some** of the actions and consequences of other users , and this will depend upon the context and tasks of the conference.

  The lack of 'controllability' (i.e. negotiable protocol for shared access to input/output devices and access to information) did cause problems occasionally during the trials, particularly in

terms of access to shared information e.g. it was impossible to modify someone else's input to the whiteboard.

- **Feedback**
  The principle of feedback to the user regarding the status and quality of the current conference call was shown to important to the users by this study.

4.3.3 Seamlessness, WYSIWIS et al.

The heuristic design principles, when applied in a limited manner are, or would be effective in the design of multi-media conferencing systems.

The limited way in which WYSIWIS is applied to the shared whiteboard was liked very much by the subjects. One area of the whiteboard design where this principle was not strictly adhered to was during page-changing. This caused great concern and lack of confidence in the system for the remainder of the session among the subjects when they realised that what they could see was not necessarily what the collaborator could see. This further illustrates the importance of the principle.

Seamlessness (Foster et al, 1988) appears, on the basis of observations, to possibly provide some answers to some of the usability problems identified during the study. The main benefit of seamlessness being that it provides what the RACE Guidance principles call 'flexibility'. This flexibility 'will allow users to adopt their own task strategies', hence less inhibited interaction with the other conferees.

5. Conclusions

The conclusions from the results, as documented in the previous section, 3 main conclusions can be drawn :

1) The multi-media conferencing system evaluated supported the conferencing task used for the purposes of the evaluation quite well, but many aspects of its design could be improved.

2) The extant HCI design principles available that were investigated (e.g. RACE Guidance (1993)) are generally valid. There is still much scope for future research into shared workspace design, and particularly into the social implications of these systems, which will play a great part in determining the uptake of these systems in the future.

3) The framework of Pinsonneault et al. (1989) for analysing the impact of GCSS (group communications support systems) on Group Processes and Outcomes is particularly useful for categorising results when studying multi-media conferencing systems. These studies could be of conferencing system usage patterns (e.g with an aim of determining user requirements for a system), or evaluative studies of the design of conferencing systems.

The study has raised many issues of concern, which are described in Section 4. From these issues, and from the literature surveyed during the course of the project, many gaps were identified in HCI knowledge. The following issues are some of the interesting and pressing unresolved questions:

- Asymmetrical use of media:

  How can users who are unable to use one type of media be supported?

- Effects of large-scale use of multi-media conferencing systems on organisations and individuals:

  How does large-scale use of a desk-top conferencing system affect organisations?

  How does long-term 'teleworking' using such a multi-media conferencing system affect individuals and organisations?

  How much face-to-face contact is required to maintain a good 'working relationship'?

- Multi-media conference metaphors:

  What are the most suitable metaphors for different conferencing tasks?

- 'Beyond being there.' (Hollan and Stornetta, 1992):

  How can video-conferencing improve on being in a face-to-face meeting?

  For example:

  How could asynchronous multi-media conferences be held to cater for parts of the world in different time-zones?

- Video quality:

  How would higher quality full-motion video images have affected the results of this experiment?

- Task analysis of meetings:

  Could task analysis techniques be used to develop a more clearly defined taxonomy of meetings and the requirements of these meetings?

- HCI theory:

  How can the existing HCI engineering design principles (Long and Dowell, 1989), such as those in RACE Guidance (1993) be extended to adequately cover design issues peculiar to 'groupware' such as social interaction?

REFERENCES

Crowley, T., Milazzo, P., Baker, E., Forsdick, H., and Tomlinson, R. (1990):
MMConf: An Infrastructure for Building Multimedia Applications. In F. Halasz (Ed) CSCW 90: Proceedings of the Conference on Computer-Supported Cooperative Work, Los Angeles. pp. 329-342. ACM.

Foster, G. and Tatar, D. (1988).:
Experiments in computer support for teamwork - Colab (Video). XeroxPARC, 1988.

Francik, E., Ehrlich Rudman S., Cooper, D. & Levine, S. (1991):
Putting Innovation to Work: MultimediaCommunication Systems. Communications of the ACM, Vol.34, No.12, pp. 53-63.

Grudin, J. (1990):
Groupware Applications: Problems and Prospects. In B. Laurel [Ed.]: The Art of Human-Computer Interface Design.,pp. 171-186. Addison-Wesley.

Handley, M.J., Kirstein, P.T. and Sasse, M.A. (1993):
Multimedia Integrated Conferencing for European Researchers (MICE): Piloting Activities and the Conference Management and Multiplexing Centre. Computer Networks and ISDN Systems, 26b(3), pp. 275-290.

Hollan, J. and Stornetta, S. (1992):
Beyond Being There. Proceedings of CHI '92,, pp. 119-125. ACM.

Ishii, H. and Miyake, N. (1991):
Toward an Open Shared Workspace: Computer and Video Fusion Approach of TeamWorkStation. Communicaions of the ACM, 34(12) pp. 37-50.

Kraut, R., Galegher, J. and Egido, C. (1988):
Relationships and Tasks in Scientific Research Collaboration. Human-Computer Interaction, 3 (1) pp. 31-58.

Long, J.B. and Dowell, J. (1989):
Conceptions of the discipline of HCI: craft, applied science and engineering. In A. Sutcliffe and L. Macauley (Eds), People and Computers V. Cambridge University Press.

Pinsonneault, A. and Kraemer, K.L. (1989):
The Impact of Technological Support on Groups: An Assessment of the Empirical Research. Decision Support Systems, 5 (2) pp. 197-216.

RACE Guidance (1993).:
Generic usability principles for service design. RACE Guidance (R1067) Deliverable 24.

Sasse, M. A., Handley, M. J. Ismail, N. I. (in press):
Coping with Complexity and Interference: Design Issues in Multmedia Conferencing Systems. in D. Rosenberg, D. & Hutchison, C. S. [Ed]: Design Issues in CSCW. Berlin: Springer.

Sasse, M. A., Bilting, U., Schulz, C-D. & Turletti, T. (1994):
Remote Seminars through Multimedia Conferencing: Experiences from the MICE project. Proc. of INET'94/JENC5, pp. 251-1/8. Reston, VA: Internet Society.

Wickens, C. D. (1992):
Engineering Psychology and Human Performance. 2nd Edition, Harper Collins Publishers.