Feasibility Study of a National High Speed Communications Network for Research and Development: Future Applications

Saul Greenberg

January, 1990

Cite as:
Greenberg, S. (1990) "Feasibility study of a national high speed communications networks for research and development: Future applications." Research Report, Learning and Collaboration Group, Advanced Computing and Engineering Department, Alberta Research Council, 6815 8 St NE, Calgary, Alberta Canada T2E 7H7. Report was prepared for Hickling Consultants.

Table of Contents

1. A Framework for Future Applications

1.1 Introduction

The proposed network will be designed to serve the needs of researchers throughout the 1990's and beyond. It is speculated that the network could begin operations at the T1 rate early in 1991. However, the network must be positioned for a smooth migration to higher speeds (ie T3 first migration). New standards and technologies will be incorporated into the network as they evolve. In fact, the network would serve as a vehicle for the development and testing of future information technology products and services. Future applications of the network must be anticipated so that the evolution of the network can be properly planned for. The information required about these future applications is the subject of this report.

This section provides a conceptual framework for future applications of the network. The framework borrows heavily from a report summarizing a National Science Foundation (USA) workshop on the *National Collaboratory*, a vision of a nation-wide structure that provides tools enabling researchers to perform their work and collaborations without regard to geographical location (Lederberg and Uncapher, 1989). The framework is a multi-tiered hierarchy (see Figure 1, adapted from Lederberg and Uncapher, 1989). At the top are the research functions in the scientific process that researchers and developers cycle through. Scientists achieve these functions through a variety of collaboration functions, which in turn are implemented via collaboration tools and underlying technologies.

1.2 The process of scientific discovery

Many of the research and development functions required by scientists map onto the steps taken in the process of scientific discovery.

1. Gain the education, motivation, and social perception of an area.
2. Build a theoretical framework and formulate questions.
3. Design an experimental approach to answer the question.
4. Assemble resources, including space, people, equipment and money.
5. Conduct experiment, possibly using project management techniques.
6. Analyze the data.
7. Form and test conclusions and theories.
8. Present preliminary results to colleagues through draft manuscripts and informal presentations.
9. Formally publish and present the results in journals, conferences, etc.
10. Obtain community reaction.
11. Establish technical applications and/or commercialize the results.
12. Attract new participants; establish a new organization; extend the project.

(Lederberg, 1989)

Research and development functions	Educate, socialize, motivate investigator	Build theoretical framework, formulate questions	Design approach		Assemble resources (space, people, equipment, money)	Project management & organization
Collaboration functions		Idea generation	Supporting community knowledge base (acquisition, assimilation & dissemination)	Collaborative resource finding Finding a collaborator	Project management & organization Capture, display, editing state of collaboration Negotiating & allocating scarce resources Decision support	Remote & collaborative experiments Telepresence Rapid remote prototyping and transport of artifacts
Collaboration tools	User training, education tools	Digital library Standards for Access	Information retrieval systems	Structured discussion	Group decision support systems	Remote experiment & resource scheduler
Enabling technologies	Software engineering Engineering technology for group interfaces Software engineering tools	Meta technologies for building required technologies & tools Instrumentation of tools to understand their impact	Standards (graphics, window systems, information...)	<u>Graphics, human-computer interface & multimedia</u>		
				Windows & virtual graphics across different machines Shared windows and workspaces	Speech/sound Video in a network & in a window	Display hardware Large public displays Portable displays & input systems

Figure 1: The collaborative framework.
Adapted from Lederberg and Uncapher, 1989
(continued on next page)

Research and development functions *(continued)*	Analyze data; Form & test conclusions, theories; Presentation to colleagues; Publish; Obtain community reaction; Commercialize; Attract new participants/form new organizations
Collaboration functions *(continued)*	Managing physical environment Service via email Ongoing/open scientific review & commentary Prepare documents Publish documents Review documents <u>Meetings</u>: Distributed seminars & meetings; Video teleconferencing; Synchronous & asynchronous communication; Shared gestures space; Informal collaboration; Social Browsing; Real time meeting management; Stuctured group history (note taking)
Collaboration tools *(continued)*	Bulletin boards Asynchronous conferencing <u>Remote meetings and editing</u>: Co-authoring tools; Shared drawing surfaces & sketching tools; Meeting room to meeting room protocols; Packet video for video meetings; Video telephones & conference rooms; Virtual hallways; Conference coordination; Cursor-passing protocols & controls for group editing
Enabling technologies *(continued)*	Robotics Tele-operation <u>Communication</u>: High bandwidth networks (gigabits?); Portable communication; Link services; Video transmission; Tools for ntetwork performance, debugging & management <u>Information databases, indexing & retrieval</u>: Document representation, editing, access; Fast linguistic search; Automated image segmentation & feature extraction; Digital image bank; Hypertext & Hypermedia; Sharable, persistent databases

Figure 1: The collaborative framework.
Adapted from Lederberg and Uncapher, 1989
(continued from previous page)

Regardless of discipline, these stages of research are a reasonable approximation of the high-level tasks performed by scientists. As the process applies to distance-separated scientists just as it does to co-located teams and individual investigators, the stages provide a meaningful motivation driving the applications the network should support .

1.3 Collaborative functions and tools

To accomplish these steps listed above, distance-separated scientists draw upon a variety of collaboration functions. They must interact casually and formally with remote colleagues; share data and resources; monitor distant instrumentation and control equipment; share information in a knowledge base; manage the project; make decisions, and so on (Figure 1, second row). Several of these functions are briefly described here, along with the collaboration tools that support these functions (Figure 1, third row); later sections of this report will expand on these.

Community knowledge base. Researchers need to acquire information from a communal knowledge base to assist their idea generation and project development. They also need to disseminate their own works into it. The community knowledge base is still mostly resident in the paper medium, leading to problems of publication delays and difficulties of finding the required information. A digital library that can overcome these problems by offering good information retrieval tools, and by allowing new information to be well-assimilated in a timely manner. Digital library should support multi-media and must set standards for access.

Finding collaborators through casual interaction. When probing for collaborators and initiating a project, researchers usually identify shared interests, argue assumptions, generate ideas, probe for suitable partnerships and extend commitments. The process often requires frequent, brief and spontaneous communications between potential collaborators. Since interaction between people is well-proven to have an exponential decay with distance (Kraut, Egido and Galegher, 1988a), technical support for casual interaction is aimed at overcoming the distance barrier.

Project management and organization. Managing a research project—especially a large ones—is a daunting task for distributed researchers. Participants require technology to assist project management. This includes not only the generation of the project plan, but also the ability to capture, display and edit the state of the collaboration. In addition, group decision support systems can assist the members to come to consensus. Technologies included project management tools, structured electronic mail tools that track researcher's commitments to each other (Winograd, 1988a), and structured argumentation tools to assist the decision-making process (Conklin, 1987).

Remote and collaborative experimentation. Remote experimentation occurs for several reasons. Because resources are scarce or expensive, they may be held by outside organizations and agencies. In this case, researchers are required to negotiate, allocate, schedule and use tools that are situated elsewhere. Alternatively, the experiment may be performed at one scientist's site, with distant collaborators being involved through remote control. Technical support of these activities falls into several areas. Robotics and tele-operation are required for remote control of the experiment. Resource schedulers assist the booking of scarce tools. Tele-data—the ability to see data at a variety of sites—is required if collaborators are to see and react to the generated experimental data. Remote prototyping is necessary if collaborators are forming the experiment interactively.

Ongoing scientific review and commentary. An important part of the scientific process is to obtain reaction from one's peers during the complete life-cycle of the experiment. For this to succeed, a researcher must be able to disseminate informally the on-going aspects of the work, and to collect the commentary back from the community. Technology can support this process in several ways. Through electronic mail, the scientist can send the information to specific people, while information can be broadcast to the community at large through a bulletin board or asynchronous conferencing system. Replies can be returned to the bulletin board, or personally to the researcher. Real-time support for distant casual interactions is another way for researchers to disseminate and gain feedback on the fly.

Collaborative document preparation. Preparing formal documents depends heavily on the styles and responsibilities of the researchers involved. Minimally, collaborators must be able to exchange drafts and comments, possibly through electronic mail. However, the future will see tools that let people interactively create and revise documents together. The technology includes real-time interactions with a document displayed simultaneously on all participants' workstations (Greenberg, 1990), or through co-authoring tools that provide structured access and annotation to a document (Leland, Fish and Kraut, 1988). The final document can be disseminated through the community knowledge base.

Meetings. The real-time meeting is perhaps the cornerstone behind the collaborative process. It is vital for brainstorming, rapid evolution of ideas and arguments, and for formalization of goals and methodology. Technology can support a variety of real-time meetings in several ways (Greenberg and Chang, 1989). Video tele-conferencing gives a sense of "tele-presence" by allowing distributed participants to see one another. Shared workspaces, along with annotation tools and a gesture space, permits participants to work together through a shared media. Meeting schedulers assist the daunting task of trying to collect all participants together at the same time. Other technologies can help a group capture and save the important points as a meeting progresses.

1.4 Outline of the report

This report structures the needs of distributed scientists into three general approaches: supporting communications between researchers (Sections 2 through 5); providing researchers access and control over remote resources and equipment (Section 6); and tools that let researchers store and manipulate information (Section 7). Although the sections do not map directly onto the framework described above, they do cover most of the collaboration functions and tools listed in Figure 1. These sections begin with an introduction to the collaborative approach provided, relating it to the specific needs of researchers and developers. Central to each section is a description of current visions, with emphasis placed upon research systems that are near commercialization or likely to be commercialized in some form over the next decade. Limitations and unknowns of the technologies are raised, and network requirements estimated. In contrast, Section 8 highlights three application areas as examples collaborative functions. Section 9 closes the report with the vision of the coherent workspace.

Computer support for real-time remote conferencing is discussed in Section 2. These are the formal or semi formal meetings that geographically-distributed people attend. Techniques for transmitting a sense of participatory presence and for sharing data over distances are covered.

Casual real-time interaction, covered in Section 3, is a crucial yet often overlooked component of scientific collaboration. While formal meetings take care of key events in the scientific process, casual interaction at the interpersonal level is the glue that keeps the collaboration working. Several innovative systems are described that bring people together through frequent, unplanned, and high-quality interactions.

Asynchronous messaging is explained in Section 4. This covers the familiar electronic mail systems we are now used to. The domain is extended to more sophisticated versions of electronic mail, including multi-media, semi-structured, and semi-formal mail systems.

Bulletin boards and asynchronous conferencing is described in Section 5. These are systems that allow people to interact with a large community by posting and viewing messages, and joining and actively participating in on-going discussions of some topic of interest. Technical variations and issues behind these systems are provided.

Access, monitoring and operation of distributed equipment is covered in Section 6. Two aspects are distinguished: access to remote computational power, and real-time systems for remote operation. Technologies such as Scada systems, master-slave tele-operation, and tele-robotics are included.

The digital library is detailed in Section 7, and lists conventional approaches such as the digital repository and bibliographic databases, as well as the attractive alternative of hypertext. The digital library is the backbone behind the communal knowledge base.

Several applications are provided in Section 8 to illustrate how the approaches of the previous sections can support high-level collaborative functions. Application areas include distance education, group decision support systems, and joint authoring. A brief note describing the need for the network for bootstrap collaborations is provided.

A vision of the *coherent workspace* closes the report in Section 9. The coherent workspace is an idealized integrated platform that supports distributed researchers. It occurs when people and their machines have a common knowledge and sense of purpose, with individual and group activities being well-coordinated.

2. Computer support for real-time remote conferencing

2.1 Introduction

Real-time remote conferencing brings participants of formal or semi-formal scheduled meetings together at the same time, even when some or all are physically distributed over different locations. This purpose is reflected in the several other names that remote conferencing goes by: multi-site tele-conferencing; distributed meetings; and same-time, different place meetings.

Remote conferencing is crucial for geographically distributed researchers and developers, with values and functions similar to face-to-face meetings. The list below indicates a few benefits.

1. Remote conferencing is an education tool where people can take interactive and participatory classes from distant experts (who can also be considered a “scarce resource” in a way similar to expensive and rare equipment).
2. Meetings are crucial for preliminary and on-going design sessions that involve brainstorming, sketching out ideas, creating outlines, and so on.
3. Meetings are needed by collaborators for on-going project development. Examples include setting goals, scheduling, formalizing the various stages of the project (design, methodology and analysis), and discussing and solving problems.
4. Draft seminars are presented to other collaborators for feedback and further development before going public.
5. Meetings are the backbone behind informal and formal research presentations to the community at large. Not only can one researcher present preliminary, on-going and final results to distributed colleagues, but the audience can provide immediate feedback, ask questions, and so on.

This section will describe computer support for real-time remote conferencing.
Two aspects of conferencing are distinguished in the following two subsections: *tele-presence,* a way of giving distributed participants in a meeting a feeling that they are in the same meeting room and the means to orchestrate the conversation, and *tele-data,* a way of having participants bring into the meeting the materials and on-going work they wish to share with others. Illustrations of several research systems furnish an image of what we can expect from tele-presence and tele-data over the next few years. In the next sub-section, an innovative and futuristic approach called *cyberspace* is described, a potentially powerful means for conferencing, data sharing and visualization that we may see in limited use by the end of the decade. The subsequent sub-section raises several limitations and unknowns that, in spite of the apparent needs for remote conferencing, hinder its use. The section closes by matching the network capabilities to the requirements of real-time remote conferencing.

2.2 Tele-presence.

Perhaps the most interesting aspect of remote conferencing has been in the field of *tele-presence*—a way of giving distributed participants a feeling that they are in the same meeting room. Tele-presence concentrates on transmitting both the explicit and subtle dynamics that occur between participants. These include body language, gestures, eye contact, meta-level communication cues, knowing who is speaking and who is listening, voice cues, and so on. Tele-presence facilitates effective management and orchestration of remote meetings by the natural and practised techniques used in face-to-face meetings.

One simple experiment in tele-presence is found in *Cantata*, a text-based remote conferencing tool that includes a multi-window broadcast environment (Chang, 1986; Chang, 1989). Participants converse with each other through text, and each participant's text is displayed through a window on the screen. An attendee has the option of selectively viewing any other participant's input to the conversation by raising that participant's window. When text is typed by one or more "speakers," it is broadcast to all participants and displayed in the window representing that speaker.

Although simple, Cantata supports tele-presence through several devices.

1. *Knowing who is speaking.* Listeners know who is "speaking" because they can see text appearing as a sender is composing it. This is especially important because, unlike voice, many people can simultaneously broadcast text.
2. *Focus of attention.* The "listener" has the option of paying less attention to specific speakers by adjusting a *focus of attention* gauge, resulting in the text being filtered to show only occasional words. The less the attention, the less actual words displayed. Listeners still get a "background hum" by seeing dots printed as the words are composed. The speaker has a corresponding gauge that indicates how much overall attention the group is paying to him.
3. *Interruptions*. A person can force everyone else to pay attention to him via a text-based equivalent of interrupting a meeting through shouting. Cantata allows any participant to compose and broadcast an interruptive message to other attendees. The sent message appears in its own window popped up on top of all other windows on the receiver's screen. Unlike the normal broadcast environment, participants cannot disallow, hide, or filter the "shouted" text.

Text-based communication, although applicable in some situations, is likely ineffective for the majority of remote real-time meetings. Most research and commercial efforts have investigated visual and audio tele-presence through *video conferencing* across dedicated meeting rooms. One or more people meet in a room; other participants in the conference meet in equivalent rooms at the distant sites. Video images of the attendees are then transmitted between these rooms across a high-bandwidth communication channel. In the simplest case, a camera will just transmit an image of all participants in the room, perhaps with the camera focusing on the active speaker. A more complex scenario would see a single screen for every participant, where monitors and speakers are all located in the same relative position across all rooms so that eye contact and directional sound cues are maintained.

Perhaps the best effort in tele-presence is MIT Media Lab's *Talking Heads* (MIT, 1983a). A remote participant is represented by a translucent mask (cast at a previous time) of his face. The video image is projected into the mask, giving the effect of a 3-d "hologram." The innovative aspect is that the mask rotates to reflect the actual head movement of the person, as picked up by motion sensors. The feeling of presence is striking. Participants can naturally and effectively orchestrate the conversation through normal eye contact and head movement.

Low bandwidth versions of Talking Heads do away with the video signal by transmitting only the head movements and the audio signal, and then selecting for display one of several pre-stored images or caricatures of the speaker that best match the incoming signal—a speech recognizer is used to match lip movements (MIT, 1983b). Although this facade is not nearly as rich as the full bandwidth version, participants still see who is speaking and are able to direct attention to each other.

2.3 Tele-data.

Most real meetings require not only the people, but also the materials and on-going work participants wish to share with others. These include notes, documents, plans and drawings, as well as some common work surface that allows each person to annotate, draw, brainstorm, record, and convey ideas during the meeting's progress. Given that an individual's work is commonly centred around a workstation, the networked computer can become a valuable medium for people to share on-line work with each other. *Tele-data* allows participants at a meeting to access, present and interact with materials that would normally be inaccessible to the distributed group.

For example, one multi-site tele-conferencing setup that uses several types of tele-data is the Multipoint Interactive Audiovisual Communication (MIAC) audiographic conferencing system (Clark, 1989). MIAC supports remote communication through transmission of high quality audio, facsimiles, still picture TV frames, real-time *tele-writing*, and chairman control of interactions over a 64 kilobit/second communication channel. Its salient features follow.

Audio. Each participant has his own microphone. Listeners receive an indication on their display of who is currently speaking. In a non-conducted meeting, anyone can speak at any time. In a conducted meeting, the chairperson can speak whenever he wishes, while other participants must explicitly request the floor from the chair.

Video. Still picture TV is used to transmit a single video frame between meeting rooms.

Writing/Sketching. Using a data tablet, participants can exchange handwritten information in real-time via the *tele-writer.* Three tele-writing scenarios are possible: exchanging the tele-writer image only; superimposing the image on the still-picture TV; and moving the cursor over the display.

Facsimiles. A facsimile can be loaded and sent from one site to another through a facsimile machine. MIAC mediates the point to multipoint communication.

Messaging. Short text messages can be sent between participants.

Another approach to tele-data stems from taking a standard computer application and sharing it between participants of a remote conference through a "shared screen" or "shared window." Each participant sees the same image of the running application on his own screen, and has opportunity to interact with it by taking turns. Special "*view-sharing*" software would allow *any* unaltered single-user application to be brought into a meeting; the application itself would have no awareness that more than one person was using it. The view-sharing software's responsibilities include registering participants, maintaining consistent shared views, managing floor control for serial input to the application, and allowing attendees to gesture and annotate around the view (Greenberg, 1990). Although simpler in idea than true multi-user applications that are aware of and cater to all participants, the capability of sharing views and interactions with the many single-user applications now available can augment significantly people's ability to work together, both in face-to-face and remote encounters.

Shared view systems are far from new. Over twenty years ago, the visionary Doug Engelbart held what was probably the first shared screen conference through his NLS system (Engelbart and English, 1968), where six displays were arranged on a table so that a group of twenty participants could see the screens. While only one participant could control the screen, other participants could control a large arrow (the first tele-pointer). Since then, shared screen systems have evolved to match current interface capabilities. *MBLINK*, for example, not only allows multiple workstations to share a screen bitmap, but also displays each participant's distinctive cursor on the view (Sarin and Greif, 1985). Several research systems now permit people to share and arrange individual windows rather than the complete screen, achieving greater flexibility by allowing one to arrange his personal display to include both private work and shared windows (Ensor, 1989; Ensor, Ahuja, Horn et al., 1988; Gust, 1989; Lantz, 1986). At the Alberta Research Council, Greenberg has decoupled the view-sharing kernel from the interface required for explicit floor control, resulting in a system that can be readily specialized to the needs of the participants and to the hardware requirements (Greenberg, 1990). On the commercial front, Farralon Software sells a simple, inexpensive but surprisingly effective shared-screen facility for the Macintosh called *Timbuktu* (Farallon, 1988). A detailed description of shared view systems is available in Greenberg (1990).

Other quite sophisticated applications of tele-data are based upon a communal multi-user workspace that allows all participants to share, view and simultaneously interact with a work surface. Perhaps the best and most well-know example is the software suite available in *CoLab* (Stefik, Foster, Bobrow et al., 1987b). Although designed by Xerox PARC to support face-to-face meetings by small groups, its software is relevant to remote conferencing. The CoLab software suite consists of three tools: *Boardnoter*, a shared chalkboard; *Cognoter*, a tool for brainstorming and idea organization; and *Argnoter*, a tool to organize and evaluate arguments.

Boardnoter supports and extends the notion of a chalkboard, and is a good example of how computer-based technology can mimic and augment the capabilities of existing media devices. Like a chalkboard, the Boardnoter supports informal freestyle sketching and erasing on a communal work area by participants. Unlike a chalkboard, the Boardnoter extends the drawing power of the group:

- text and figures to be easily movable;
- images to be re-organized on the display and stored in databases;
- active involvement of participants from their seats;
- multi-person gesturing via a tele-pointer (Stefik, Bobrow, Foster et al., 1987a).

Meeting tools need not be as generic as a chalkboard. Using existing equipment, a meeting can bring in a particular flavour of meeting tool designed to facilitate a special meeting task. Cognoter is an example of such a tool, for it is explicitly designed to separate the planning of a presentation into three separate stages (Foster and Stefik, 1986; Stefik et al., 1987b). In the first stage, ideas generated by each participant are entered simultaneously as independent "catch-words" in the free space on a public window. Catchwords can be further annotated with supporting text in sub-windows. In the second, ideas are organized by allowing participants to order, link and cluster the displayed catchwords. Physical and conceptual clutter is reduced by allowing people to chunk grouped items into a hierarchy. The third stage involves idea evaluation. The overall structure is reviewed, details added, and irrelevant ideas removed.

2.4 Cyberspace.

The most innovative and futuristic approach to remote conferencing may lie in *Cyberspace*. Cyberspace immerses a person's senses into a three-dimensional (3-d) simulated virtual world. Seeing the world in a stereoscopic head-mounted display that has a screen for each eye, one moves through it by head and body gestures. Motion sensors pick up and translate real movements to virtual ones, and the view is adjusted accordingly. Users interact with the simulated world through a data-glove or data-suit that allows them to grasp and manipulate the virtual objects they see. They hear sounds through a 3-d audio display. The effect, although still primitive, is to exist and interact within a virtual reality—cyberspace.

The relevance of cyberspace to remote conferencing becomes apparent when two or more people interact within the virtual space. Imagine a conference held in virtual room, with attendees milling about, holding private conversations, and viewing and manipulating some of the available 3-d data entities. Science fiction? Not quite. The first demonstration of VPL Research Inc's *shared virtual reality* system occurred on June 7, 1989 in San Francisco[1].

The potential of cyberspace is far beyond anything offered by other remote conferencing systems. Not only can participants interact with each other (eg by moving around a virtual room), but the meeting place can simulate any environment. For example, the cyberspace model may be constructed and animated from incoming data points of an instrumented experiment (or the simulation of an experiment), with scientists viewing the data and adjusting the control parameters interactively. Alternatively, the model may be one of a conventional meeting room with a sophisticated whiteboard, so that participants can act within a cyberspace meeting as they would in a real meeting.

2.5 Limitations and unknowns.

Although both tele-presence and tele-data are clearly important, their effective implementation and the role they play in particular types of meetings are not well known. Requirements of meetings vary greatly; group dynamics are volatile; people are not used to computers in meetings; software for meeting support is (at best) at the prototype level; technology is obtrusive.

Video conferencing, for example, has fallen far short of its promise for several reasons.

1. Vendors gave video conferencing an ill-conceived image as a means of reducing the need for travel to face-to-face meetings, which it does not do (Egido, 1988). Travel actually increases, for the need for direct meetings grows with the frequency of the interpersonal contacts made over video.
2. Although video conferencing has proven suitable for passive meetings emphasising presentations, it appears to be a poor medium for supporting the more common highly interactive style of meeting where there is much inter-personal interaction (Egido, 1988).
3. Video presence of participants seems to add little to communication (Chapanis, 1975; Johansen and Bullen, 1984), and therefore may not warrant the technological and physical restrictions it places on the meeting.

---

[1]Off the shelf equipment to realize 3-d cyberspace is available through VPL Research Inc, Redwood City, California.

Anecdotal evidence supports these views. One account mentions that after a six-month novelty period had worn off, the day-to-day use of the video conferencing facility shifted from display of participants to simply pointing the camera to the data. Participants were content to talk anonymously over a speaker phone (Guttman, 1989). It is the data that was considered more important than the simple form of tele-presence offered by video.

Some of the failures above may not be due to the notion of tele-presence, but rather to limitations in technology. Special video conferencing rooms, for example, means that participants must schedule and limit their meetings to these rooms rather than use their own offices. (Although cyberspace conferencing may remove these restrictions, the technology is costly and the effects still primitive.) The high cost of bandwidth means that video transmission often uses a compression scheme that severely impacts on the quality of the displayed image. Furthermore, keyboards, monitors, microphones, wires, and cameras may be a significant intrusion to the meeting, particularly to those participants who do not feel comfortable with the technology.

Given that face-to-face meetings are considered more effective than remote ones, there is still much room for improvement. Some progress is being made. Xerox PARC, for example, uses multi-disciplinary teams of sociologists, anthropologists and computer scientists to study how people communicate through a video channel (Harrison, 1989; Stults, 1988). The knowledge acquired through this process will help determine the design requirements behind truly useful remote conferencing systems.

These limitations should not discourage use of remote conferencing. Technical problems will be overcome and the human factors issues solved. The real potential exists because remote conferencing is vital to overcoming the distance barrier faced by distributed researchers and developers.

2.6 Network requirements

A 1.5 Mbps network can address some of the requirements of remote conferencing. In particular, the tele-data requirements that have already been developed and used on LAN networks can migrate to the 1.5Mbps network with little performance penalty. These include screen sharing of existing single user application packages, and multi-user workspaces that allow participants to simultaneously draw and type in the view.

The network may also suffice for some limited types of tele-presence. Textual real-time communication systems such as Cantata require modest bandwidth, while the shared gesturing ability found in workspaces probably requires little more than the transmission of a person’s pointer location. As mentioned above, low bandwidth versions of Talking Heads does not require video transmission (although there is still a question about how well a voice-animated image provides tele-presence). Similarly, slow-frame video can capture at least some aspects of tele-presence.

Yet the dynamics between attendees in a meeting can be so diverse that it is unlikely that any limited channel will capture all the subtle yet important cues. It stands to reason that real tele-presence would require at least a full video channel and a corresponding increase in bandwidth over the network. At the very least, one video connection is needed for room to room contact. Bandwidth demands multiply with the connections made, as seen when the number of rooms increase, or when video screens represent each individual in the meeting. A similar high bandwidth requirement is needed when tele-data requires video, as may occur when the data is not available in digitized form, is not easily scanned in, or is animated.

The actual network requirements for cyberspace is unclear. Three-dimensional real-time animations must be modelled and exchanged; interactions between participants with each other and with the simulated environment transmitted and synchronized; and sound (including its coordinates of origin) transmitted. Although simple cyberspace interactions using distributed models may run over the 1.5Mbps network, a true realization of an alternate reality would likely require considerable bandwidth.

3. Casual Real-time Interaction

3.1 Introduction: Patterns and requirements for scientific research collaboration.

Real-time conferencing addresses scheduled formal and semi-formal meetings. Yet it is not necessarily pre-planned, purposeful meetings that are best supported through computer mediation, but casual unplanned meetings as well. This section argues that real-time support for casual interaction is vital to scientific research collaborations and must be supported by the network. Several innovative systems indicate the potential shape of these future applications.

Research is fundamentally a social process, with many interactions required for people to initiate and execute collaborative research. Yet only a small part of this process relies on the formal communications supplied through scheduled meetings. Consider the model of research collaborations below, developed from interviews with scientists in various areas of research (Kraut, Galegher and Egido, 1988b).

	Initiation	**Execution**	**Public presentation**
Relationship level	Finding a partner Sharing background assumptions Identify shared interests	Supervising and sustaining progress Establishing division of labour Establishing trust	Establishing division of credit
Task level	Generating ideas and planning	Sharing information Coordinating activities Doing the work	Writing the manuscript

The process of many of these stages clearly rely on frequent casual contacts to maintain collaboration on both the task and the interpersonal level. In initiation of a research project, for example, researchers usually identify shared interests, generate ideas, share assumptions, probe for suitable partnerships and extend commitments. During execution of the project, collaborators must share information and coordinate activities, reshape project goals, establish sub-goals and divisions of labour, comment on the other's work, and so on. Although scheduled meetings may be held during this process, most communication is informal. Through brief communications occurring in quite casual settings—over coffee or lunch, in hallways, or perhaps during an otherwise uninteresting meeting—collaborators exchange ideas and information, feedback, encouragement, status changes and minor (yet important) points. The technologies used are simple—a scrap of paper to make a point or jot down ideas, or perhaps a whiteboard for brainstorming. Kraut notes that collaborators rarely used formal project management techniques to assist their task (Kraut et al., 1988b) .

Kraut concludes that the establishment and maintenance of a personal relationship is just as vital to the collaborative research effort as the content of the work itself. Three criteria are usually necessary for informal communication and the resulting collaborations to succeed (Root, 1988):

1. *High frequency of communication*, where potential collaborators have ample opportunity to communicate to each other;
2. *High-quality real-time interactions* allowing intense face-to-face sessions and the sharing of materials;
3. *Low cost interactions* so that potential collaborators can have spontaneous or one-person initiated meetings with little effort.

Root especially emphasises the importance of brief unplanned encounters where bits of technical and personal information are exchanged "on the fly."

Given these criteria, it should be no surprise that the bottleneck to rich spontaneous interactions is distance. Frequency of communication decreases, communication media is often low bandwidth and of low quality, and interactions come at a high cost (eg telephone tag or travel time). Interaction between people is well-proven to have an exponential decay with distance, as shown not only by the reduction in the frequency of face-to-face interactions (Allen, 1977; Kraut et al., 1988a), but also by the lower exchange of telephone calls (Mayer, 1977) and electronic mail messages (Eveland and Bikson, 1986). For example, the number of collaborations drops off sharply when one contrasts people working on the same floor, on different floors, and in different buildings. Further study indicated that the effects of physical proximity and interaction are not merely an artifact of people with similar interests usually being co-located (Kraut et al., 1988a).

The conclusion formed is that distributed researchers are less likely to go through the critical initiation process if they cannot get connected, leading to a decrease in potential collaborations. Similarly, the communication problems encountered during the execution and presentation stage may be significant enough to hinder or abort the project.

Technology has potential to bring distance-separated people into contact through frequent, unplanned, high-quality, and real-time interactions that come at low personal cost. Two visions are described in the sub-sections below: *video hallways* and *shared alternate reality*. Both require high-bandwidth communication networks to succeed. Limitations and bandwidths are raised in later sub-sections.

3.2 Video hallways.

Several research laboratories are exploring the possibility of "video hallways" for casual interaction between remote sites. The first case was Xerox's *Video Wall*, which placed a slow-scan video connection between two research laboratories located in California and Oregon (Goodman and Abel, 1987; Stults, 1988), also summarized by (Root, 1988). Spontaneous "drop-in" interactions between people at the two sites were encouraged by placing large video screens in common areas. Point to point connections between individual offices were also allowed to a limited extent.

Video Wall worked. Goodman and Abel reported that 70% of all Video Wall interactions were spontaneous, and the other 30% planned (Goodman and Abel, 1987). A different breakdown indicated that one third of all communication was social in nature and two thirds technical. Users reported that these interactions would probably not have taken place without the link. On the other hand, users reported dissatisfaction with the poor image quality of slow scan video.

A second prototype video hallway is *Cruiser* (Fish, 1989; Root, 1988). While Xerox's Video Wall directly connects two physical locations, Cruiser attempts to create a virtual community where everyone has instant access to everyone else. Cruiser is designed on two premises: 1) users can browse a virtual world seeking social encounters, and 2) users can construct, organize and populate the virtual world independent (within reason) of the physical world. There are three methods for browsing.

1. A *jump* supports a direct planned movement to a physical location. A user can select a specific location, and the image projected from the camera at that location appears on his screen.
2. A *path* extends the jump idea by listing a sequence of locations and the order in which to visit them. This, in effect, becomes a "virtual hallway" through which the user can walk through.
3. A *random walk* is similar to a path, except that the Cruiser system generates the locations in the sequence. The selected locations can be purely random, or they can be chosen as a function satisfying some user desire.

What do people actually see when using Cruiser? A visitor peeking into a person's office (via a video camera) will see and hear whatever image and sound the camera projects. The occupant, on the other hand, sees a virtual hallway on his screen and the image of the visitor as he is passing through.

While "peeking" into offices raises the spectre of George Orwell's "Big Brother," the Cruiser design recognizes an individual's desire to control privacy. As in real-world offices, people have the option of metaphorically keeping the door open (seen by the visitor as full video); semi-opened (seen by visitor as partially-drawn blinds across the image); or completely closed (no image projected). Root identifies several variables that can be controlled for setting privacy levels: access of visitors to the video and audio channel; ability of visitors to interrupt; ability of selected visitors to over-ride other settings; and a privacy value (Root, 1988). For example, a closed door policy is implemented by setting video and audio transmission to none but allowing interruptions by people with high priority levels.

Many interactions are not a result of someone "cruising for action." Rather, they arise from people bumping into one another while performing their everyday work, and in joining in on conversations already in progress. Cruiser supports this style of "situated interaction" by allowing people to attach images to several work activities. For example, sending a document to the printer will automatically invoke a random walk that encounters the other people using the printer.

Root warns in his paper that Cruiser is very much a prototype, and much remains to be built and tested. In contrast, the *US West* TeleCollaboration project supports a high-speed full video link between two research laboratories in Colorado separated by a distance of 100 kilometres (Corey, Abel, Bulick et al., 1989) . Their system allows users to "video walk" through the physical hallways and offices of distant sites. They can even scan larger rooms by tele-operation of remote-controlled video cameras. On their own screen, they can search for people, see who is around, start informal conversations, engage others in coffee room chit-chat, and so on.

3.3 Shared Alternate Reality.

Video hallways use computer support only to help establish personal encounters. Its users have to decide to go cruising, or they have to leave themselves open to intentional or accidental encounters. Randall Smith's vision of the *Shared Alternate Reality Kit* (SharedArk) takes another approach (Smith, 1988). Unlike virtual hallways, SharedArk is more than a medium for meeting people.

SharedArk is based on a model of a shared virtual yet physical world—a two-dimensional "flatland"—used for teaching students physics. Students can wander through flatland, and manipulate physical objects with a mouse-operated hand. Unlike most virtual worlds, flatland is populated by all the people travelling in it. Students may accidentally encounter each other (ie one will see another person's hand). They then have opportunity to open an auxiliary video and audio connection for more direct communication. Within SharedArk, students can form collaborations on simulated physics experiments and jointly edit text and graphics.

SharedArk has several other features. First, a person can see who and what is around him through a "radar view" that provides a miniature of the surrounding space. Second, people (and their objects) can quickly move from one virtual site to another by stepping into a "tele-porter." Third, people can set up private regions within the virtual world that excludes other people from travelling or looking within it.

A natural extension to SharedArk is, of course, cyberspace. Pushing Smith's vision a bit further, imagine a multi-world version of cyberspace, where each "world" represents a topic of interest. People travelling through and exploring a particular world of their choice have opportunity to accidentally encounter and make acquaintance with other people with similar interests, just as travellers do in real life.

3.4 Limitations and unknowns.

While video hallways are exciting, they have several serious limitations. First, the cost alone is prohibitive. High-bandwith video between distant sites incurs expensive communication tariffs, and low-bandwidth video is only marginally acceptable to participants. When this is combined with the cost of cabling and adding cameras and monitors to all offices, it is likely that most institutions would consider the simpler video wall setup before video hallways, at least in the short term. A second limitation is social. Although many people are comfortable with working in common work areas or within open offices, video "peeking" raises the concern of invasion of privacy and abuse by management. A third limitation is one of effectiveness. While informal interaction is important and does seem to work for small groups, the possible influx of visitors to an office may lead to people—especially popular people—to adapt a closed door policy. Distance, after all, does provide a buffer against excessive drop-in visitors.

Video hallways have explored personal encounters only. If these meetings could be augmented by importing work tools as well (eg drawing surfaces), we may see the start of the "back of the envelope" computer-supported spontaneous meeting that allows people to develop, jot down and share ideas in a non-formal situation. SharedArk is, of course, closer to this data-oriented model. Smith, however, points out severe performance problems in his system, mostly due to the inadequacies of present-day network and computational technology for handling real-time distributed interaction and animation. In practice, SharedArk will only work well with three or four users and a few dozen objects (Smith, 1988).

Another aspect that must be considered in casual interaction is the “startup problem.” Random encounters, while potentially rewarding, can be wasteful of effort, particularly when one person actively seeks but cannot find other specific people to meet with. The task of establishing a casual encounter with a specific group is even more frustrating than scheduling a meeting. Not only must each person be contacted individually, but it is virtually certain that not everyone is available then and there. Even collaborating on a casual basis with those who are available electronically (ie actively working on their terminals) is difficult without mechanisms for easily identifying that availability. For example, experience with Cantata has shown that people tend not to use it because they have to phone or walk down the hall to get others to activate the program; the cost for a casual remote meeting is just too great.

The problem is that people connected on networks through workstations have few ways of knowing (without video hallways) who is present and available for conferencing, even if they are signed on. Two mechanisms are being implemented by Chang and Copping at the Alberta Research Council. The first, *Messenger*, is a Macintosh desk accessory that shows a person all others who have recently moved a mouse or touched a key within the last few moments (currently set at 60 seconds). A message can then be sent to a subset or all those identified as recently active. Those receiving the message hear a bell and see a flashing icon and can then view the message and reply. The second, *Golf Ball*, is a common knowledge mechanism based on the *Messenger* kernel. A message containing a proposed action (eg calling a meeting) is sent to a group of recently active persons, and the message acceptance and acknowledgments and responses to the proposal are seen by all who accept the message. Thus the entire group is aware of the responses of the members as they accumulate. The success of these mechanisms in promoting casual group formation in networks will be seen with experience.

3.5 Network requirements

A prerequisite of casual interaction is that sessions must be of high-quality and real-time. By definition, then, the network requirements will be high. At the very least, there must be enough bandwidth on the network to support: a method for allowing potential collaborators to know who is around; a voice channel; a shared workspace (tele-data). These demands can be met within the constraints of a 1.5 Mbps network.

Truly effective interaction likely requires a video link. As with remote conferencing, slow-scan video could work at the lower bandwidths, while compressions schemes may allow full bandwidth video across a 1.5 Mbps network. This allows, however, for only a single connection. While this is reasonable for a system like Video Wall, both Cruiser and the TeleCollaboration project would require multiple video links. The alternative is to allow only one person to monopolize the network while wandering the virtual hallways and while establishing a link. A similar problem occurs within SharedArk. While initial contact is made through the multi-user “virtual world,” the ensuing interaction establishes a video link.

4. Asynchronous Messaging

4.1 Introduction

Perhaps the greatest success story in bringing people together over a network is in computer support for asynchronous communication—non real-time discourse and information exchange between people. This section will concentrate on one aspect of asynchronous communication: electronic mail and its variations.

Electronic mail (email) is an important support tool for geographically-distributed and co-located researchers and developers. It serves a role far beyond standard surface mail, extending into roles traditionally held by inter-office notes and memos, the telephone, and real-time meetings. Although benefits overlap, email affords different capabilities.

- It addresses communication where the sender does not require immediate response. Similarly, it is well suited for those cases where the recipient needs time to formulate an answer.
- It is usually one-person initiated.
- It overcomes “telephone-tag.”
- Delivery time is significantly faster than surface mail, with message transmission in the order of minutes and even seconds being common. In this way, email should be considered more of a messaging system rather than a mail system.
- It comes at low personal cost. Electronic mail is usually quite informal in content and format, and fairly easy to compose and read.
- It is a means for “principals” to communicate with each other. Intermediaries (such as secretaries) do not usually compose electronic mail as they would with written mail.
- It crosses the boundaries between organizational and political hierarchies. Partially due to the lack of formality in message contents, the absence of a person’s organizational title in a message address and the low personal cost of sending a message, junior people are more aggressive at communicating with senior people. This is in marked contrast to the way other communication devices such as telephones, face-to-face contacts, and surface mail are used.

The points below indicate several benefits of email to researchers and developers in a distributed community. This list can be extended to almost all aspects of communication.

1. *Establishing contact with other researchers.* One person may initiate contact with another person, even though no previous direct communication had taken place. The contact may be motivate by several reasons, such as when a reader of an article wishes to pass comments on to an author, or when one is aware of another’s common interests.
2. *Supporting communication in all stages of a research project.* Virtually all stages of research demands communication between collaborators, from project initiation, design, management, analysis and dissemination.
3. *Acquiring information.* Researchers may request information directly from the most appropriate person. For example, acquisition for a rare resource may be preceded by a request asking about its availability, cost, and so on.

4. *Dissemination of research.* Researchers may post early draft versions of papers to colleagues and collect back comments. Similarly, developers may send out product information, press releases, and announcements to specific parties.
5. *Dissemination of information to a known group of people.* Since mail also supports one-to-many communication, people may post items of interest to members of a distribution list.
6. *Document transfer.* Email is not restricted to personal messages. It may also be employed as a simple way to transfer documents or other digital resources.
7. *Social contacts.* As mentioned in the previous section, continual social contact is the glue that helps collaborations persist. Email is a reasonable way for distributed people to exchange social notes, keep people updated, and so on.

The section begins with a background of the basic electronic mail service service. Users compose and read textual mail through any one of several user interfaces, while the mail itself is conveyed through a separate sub-system that routes the message over various networks until the desired destination is reached. I detail why such a simple service has proven effective, and what general hurdles are still in place. The section continues with a description of several variations on email that go beyond simple unstructured text messages. Through *multi-media mail*, people can exchange almost any digital information over the wire—graphics, animations, voice, and so on. Through *semi-structured messages,* people can compose typed messages through form filling, and have the system filter incoming mail through its knowledge of the message structure. Finally, *semi-formal messages* restrict the inter-personal exchange of messages to a pre-defined protocol.

4.2 Basic electronic mail

4.2.1 Overview. Through a basic electronic mail service, people can compose a text message with a text editor and send that message to one or more other people. This is the service that most of us are familiar with. The software (and the user interface) usually distinguishes a message into a header and a body. The header normally contains routing information such as who sent it, who will receive it, the address, time of mailing, whether the note has been carbon copied, forwarded, replied to, and so on. A subject line summarizing the note is usually included in the header. The body, on the other hand, contains the actual text message being mailed and is usually just passed through as is to the receiver.

Electronic mail has proven a surprisingly effective means for asynchronous communication. Its advantage over surface mail is its speed. Despite even large distances, turn-around time can be seconds or minutes. Because of this swiftness, email is closer to a "messaging" system, and often augments the roles more conventionally assumed by inter-office memos and even telephone calls. Email is especially good at getting around the time-wasting "telephone tag" that haunts attempts at real-time conversation. Sproull and Kiesler summarize this view of email.

> Most organizational analyses of electronic mail view it simply as an information accelerator, a tool that reduces the amount of time it takes for people to get information they otherwise would have received more slowly
> (Sproull and Kiesler, 1986).

There are several other points worth making.

- Email correspondence is typically informal, especially when compared to surface mail. Short, mis-spelled and grammatically incorrect email messages are common.
- The availability of distribution lists means that a user can trivially send the same message to many people.

Studies of electronic mail within organizations have verified its value. For example, consider the following case study where two slightly different groups were observed in a natural office setting (Eveland and Bikson, 1988). Both groups comprised two types of people: normal employees who worked in the office, and ex-employees (retirees) who were usually at home. While members within each group could communicate between each other in conventional ways, one group also had basic email facilities. They found that retirees using email had a much higher rate of communication with other members of the group when contrasted with those who did not have email (the control group). Also, communication did not cluster as much when email was used—the boundary between cliques was not as sharp. In this case, email was successful in keeping retirees and their valuable expertise involved with the office. In contrast, the office eventually lost contact with retirees in the control group. The overall interpretation is that email significantly and directly affects the outcomes and the process of cooperative work.

4.2.2 Limitations and unknowns. Electronic mail use is not without its limitations. Several technical problems are described below (see (Pliskin, 1989) for added detail).

- There is no good general directory service for querying a person’s address, and addresses are often unstable.
- Addresses are rarely independent of a network or particular mail routing service. Even when an address is found, a person may not have access to the transport network specified by the address.
- Email is often only as reliable as the nodes in the network. When (say) a recipient’s computer has been down for even a moderate amount of time, the mail may be returned to the sender.
- Interfaces to email systems range from good to terrible. In some interfaces, sending a quick note can be a tedious process, and sorting and filtering incoming mail a difficult chore.

Socio-emotive problems and issues mostly stem from an individual’s acceptance of email, having a critical mass of people using email, and relying on people to log onto the system on a regular basis.

- The use or non-use of email by key people in the organization sends a clear signal to subordinates on how they should use it.
- Some people do not like electronic mail. They may have prejudices against working on a computer in general, or they consider the overhead of learning an operating system, a mail sub-system, and an editor not worth the bother.
- If people do not have ready access to a terminal, the time to receive a message will be longer. As a corollary, email is used most heavily by people with a terminal on their desk (Sproull and Kiesler, 1986).
- As electronic mail is not delegated as easily as paper mail (ie to a secretary), senior managers may be deterred from using it.
- Junk mail (even well-informed junk mail) is beginning to inundate email systems. Unlike physical junk mail, the recipient often has to read part of the note before recognizing it as junk—there are no packaging cues.

- Email usually lacks social context cues. Little is sent about an individual, such as one's position within the organization, job title, age, and appearance. Email also lacks cues that indicates the situation in which the mail was created and how it should be received (eg corporate memos, announcements) (Sproull and Kiesler, 1986).

Sproull and Kiesler (1986) noticed several interesting social effects from their study of email use.

- People preferred to send email messages to superiors than to subordinates.
- People preferred email for sending bad news.
- People behaved irresponsibly more often with email than with face-to-face conversations.

Electronic mail is more than just a communication system, for it also demands a variety of time and task management activities from its users. Mackay studied email users and rated them in several categories, each with quite different habits and objectives (Mackay, 1988). Prioritizers concentrate on the problem of managing incoming messages. Archivers concentrate on archiving information for subsequent use, and delegators delegate mail by passing it on to others. Mackay's study indicates that mail use is strikingly diverse, and that designers of email interfaces should recognize this diversity by creating systems that provide flexibility over a wide range of users.

4.3 Multi-media Mail

4.3.1 Overview. Basic email achieves a high degree of success because its messages presume a simple yet ubiquitous standard—ASCII text. Since most computers are well equipped to handle and manipulate ASCII, the person has a rich and usually familiar environment in which to compose, edit, and save mail. While email does not demand this standard—any binary stream can be sent in the body of a message—the interface is rarely set up to interpret the stream as anything but text. It is up to the recipient to recognize unusual message bodies and to view it through the appropriate external software.

Of course, other non-textual standards exist. Facsimile machines, for example, scan and transmit documents as images. The advantage is that any black and white document can be sent as a facsimile. The disadvantage is that the units of the message—characters, drawing components—are no longer available for manipulation by the receiver.

Multi-media mail, on the other hand, is built on the premise that mail is more than just text. Mail systems should allow many different types of media objects (or combination thereof) to be composed, sent and received as messages (Borenstein, Everhart, Rosenberg et al., 1988). A multi-media message, for example, may comprise a mixture of plain text, a drawing, a voice annotation, a video. The receiver not only sees the message in its proper form, but also has the ability to manipulate the different media objects through integrated on-line tools.

Perhaps the best example of multi-media mail is the Carnegie-Mellon University *Andrew Messaging System* (Borenstein et al., 1988; Borenstein and Thyberg, 1988). Its important points are that it is a combined mail/bulletin board facility, and that it is multi-media. One can, for example:

- transmit line drawings, rasters, animations, and spreadsheets;

- ask for responses to a message via mail that asks its reader to select from a list of choices;

Interestingly enough, the multi-media aspects of the Andrew Messaging System did not catch on immediately with its user population. Instead, Borenstein and Thyberg (1988) describe a gradual “raising of consciousness” over several months. The major breakthrough came when a user mailed a bug report by including a screen snapshot of the clearly visible bug instead of a verbal description. Through the examples of how Andrew was actually used over a long period of time at Carnegie Mellon University, Borenstein and Thyberg (1988) leave a positive impression of what advanced multi-media technology can offer.

Diamond is another example of multi-media messaging system that allows users to create, edit, transmit, and manage multimedia documents (Thomas, Forsdick, Crowley et al., 1985). The emphasis on documents instead of messages reflects the designers’ belief that messaging is only one part of the document handling process when using electronic mail. As with an Andrew message, a single *Diamond* document might include various media forms: structured text and object-based graphics, bitmaps, voice, and even active spreadsheets. *Diamond* also maintains its documents in a distributed database. As a consequence, users do not store the actual documents, but rather a *citation* (ie a pointer) to it. This approach means that multiple copies of documents need not be stored, a considerable savings in storage requirements given the potential size of multimedia documents.

Digitized speech and voice-messaging systems are also becoming readily available. IBM has had the speech filing system (commercialized as the Audio Distribution System) available since 1981 (Gould and Boies, 1983; Gould and Boies, 1984). Technical workstations, such as Sun Microsystems and the NeXT machine, have or are beginning to include voice mail within the standard mail system.

4.3.2 Limitations and unknowns. Multi-media messaging is fraught with severe difficulties and unknowns. Several problems are listed below.

- Aside from ASCII text and facsimiles, there is no agreed-upon standard for sending multi-media.
- Multi-media often requires high bandwidth, which is expensive in terms of both transmission and storage.
- Few workstations can handle “esoteric” media forms, such as sound, video, and color. Media requiring raster capabilities are similarly affected. Although bitmapped displays are common, they are still heavily outnumbered by character-based displays.
- Multi-media will require senders (and receivers) to know how to use potentially complex multi-media editors.

4.4 Semi-structured and semi-formal mail

Basic electronic mail has very little explicit structure beyond separating the address field from the body of the message. Yet many messages (and sequences of messages) implicitly contain some type of structure. Personal notes, memos, standard forms, meeting announcements, and equipment requests are some examples of different semi-structured messages, each containing a different structure. Proposals/rebuttals, requests/confirmations, and questions/answers are examples of semi-formal messages that follow different between-message protocols. If the structure and protocol is made explicit and knowable to the computer, there may be potential for supplying better email systems.

4.4.1 Semi-structured messages. Malone et al defines a semi-structured message as

> …messages of identifiable types, with each type containing a known set of fields, but with some of the fields containing unstructured text or other information
> (Malone, Grant, Lai et al., 1987).

He gives as an example a seminar message, where structured fields would include the seminar time, place, speaker, and topic. Some of these fields may be typed to allow only permissible values. For example, the place slot may only allow a room number from some enumerated set of possible rooms. The seminar abstract would be an example of an unstructured text field. Another example of a semi-structured message is a "request for action" that includes a deadline field.

Several advantages occur when structure is made explicit (Lai and Malone, 1988; Malone et al., 1987; Malone, Grant and Turbak, 1986).

- Computers can process structured fields more easily than free text fields. For example, unread notices about seminars that have already taken place may be removed automatically, and messages passing some specified criteria may be placed in different folders.
- The system may fill in default values or offer a list of choices to the composer of the message. For example, the "place" slot in the seminar message may default to the most commonly used meeting room, or may have a pop-up menu attached to the field allowing the person to select from a choice of rooms.
- Semi-structured mail reflects how people naturally generate, process and categorize routine messages.
- Templates encourage a consistent (and perhaps organization-wide) approach to sending particular message types. For example, if a request for action template includes a deadline slot, the composer will be more likely to use it than if he was composing a completely unstructured message.

One example of a semi-structured messaging system is Object Lens (Lai and Malone, 1988), a second-generation version of Malone's fairly well-known Information Lens (Malone et al., 1986). Object Lens contains two fundamental ideas. First, passive information can be represented as *semi-structured objects*, where each object is defined as part of an inheritance hierarchy. Consider the added structure as one specializes through the following hierarchy branch:

"Thing — Message — Action Request — Meeting Proposal"

Whereas a completely unstructured message may be a primitive mail form, the meeting proposal includes specific information about time, place, decision requests, and so on. By defining and modifying templates for these objects, users can represent and interact with many different kinds of information.

In the second idea, Object Lens profits from the added structure through *semi-autonomous agents* and *active rules* for processing information. When creating these agents, users specify rules for automatically processing information in different situations. A rule triggered by incoming within-organization mail may, for example, sort messages from superiors into an “Urgent” folder, and discard those messages arriving from an uninteresting distribution list. With these two ideas, Object Lens integrates object-oriented databases, hypertext, and electronic messaging with intelligent routing.

4.4.2 Semi-formal messages. Structure not only exists within a message, but also between messages. A semi-formal message is:

> a message exchange following a given protocol, where the message type and between-message sequences are restricted by the protocol.

A well-known and well-documented semi-formal messaging system is Winograd’s Coordinator (Winograd, 1988a; Winograd, 1988b; Winograd and Flores, 1986). The protocol used by the Coordinator is based upon speech act theory that asserts

> …every utterance falls into a small number of categories and an act within one of these categories does not occur at random. If a person requests another person to perform a task then there is only a limited set of valid responses to this request.
>
> (Bair and Gale, 1988)

There are five things you can do with an utterance (Winograd, 1988a).

Assertive. Commit the speaker to the truth of what is being said.

Directive. Get the hearer to do something.

Commissive. Commit the speaker to some course of action.

Declaration. Declare some correspondence between the propositional contents of the speech act and reality.

Expressive. Express a psychological state about some state of affairs.

Winograd goes on to describe a structure showing how speech acts relate to each other in a larger “conversation for action” scheme that represents how one person requests another person to do something. Drawn as a state transition network, the structure defines the possible interactions between utterances, where utterances are typed by request, promise, counter-offer, declination, cancellation, or report.

Winograd’s Coordinator implements the “conversation for action” protocol mentioned above on an IBM PC-based electronic mail system. Because the system is aware of the protocol, it can help a user keep track of where things stand (Winograd, 1988b). The summary screen displays the new conversations that have started, the on-going conversations one is involved in, what commitments have been made, what responses are expected by other people, and matters that have been completed. Depending on where one is within the protocol, a list of choices of allowed message types is offered to the person who wishes to respond to a message. Winograd summarizes the appeal of the Coordinator.

> The Coordinator has no magic to coerce people to come through with what they promise, but it provides a straight-forward structure in which they can review the status of their commitments, alter commitments they are no longer in a position to fulfil, anticipate coming breakdowns, make new commitments to take care of breakdowns and opportunities appearing in their conversations, and generally be clear (with themselves and others) about the state of their work.
>
> (Winograd, 1988b)

4.4.3 Limitations and unknowns. Although semi-formal messaging has not been around for very long, early usage studies are encouraging.

Semi-formal communication as exemplified by the Coordinator was studied by Bair and Gale (1988). The field study includes six companies that had installed the Coordinator on seven to 900 nodes. Their anecdotal findings are summarized below.

- Some (but not all) people, especially managers, found the Coordinator's ability to track commitments to be a powerful management tool, particularly when subordinates were geographically widespread.
- Structured conversation reduced junk mail.
- The Coordinator encouraged reflection and thoughtfulness in management and thinking.
- The Coordinator was an aid to planning through its listing of outstanding commitments.
- The Coordinator enabled formal negotiations to proceed smoothly, even when participants had no face-to-face relationship.
- A minority of companies/people reported commitment tracking to be socially unacceptable.
- Negative reactions to the Coordinator usually came from people used to a different email system and interface or to those who did not understand the theory and paradigm behind the system.
- A critical mass of users was necessary before the Coordinator became beneficial.
- Since not all "speech acts" occurred over the Coordinator (eg face-to-face and telephone meetings are other communication channels used), the tracking process was sometimes incomplete.
- The strongest predictor of success was the use of the Coordinator by upper management (Bair and Gale, 1988)

Although the comments above seem mostly positive, the Coordinator also has the dubious honor of being known as an example of "fascist software" that dictates that users do things a certain way. Since the Coordinator fundamentally alters the way people work, Bair and Gale suggest that its introduction to an organization must be accompanied by workshops that allows employees to learn the ideas behind conversation for action and to buy into the approach.

4.5 Network requirements

Most basic electronic mail requirements are well-handled by a fairly low-bandwidth network. This is usually due to the small size of the mail packet, and because the time to transmit is not critical due to the asynchronous aspect of the communication. As mail becomes more sophisticated, as with the addition of structure and protocol into the message, the average message size will rise with the increase in transmitted information.

More bandwidth is required as mail contents become increasingly sophisticated. Taking document preparation systems as a model, we have seen a progression from simple ASCII text to quite sophisticated typeset documents including complex figures and bit-mapped images. There is every reason to believe that users will demand electronic mail to follow the same course.

The real increase in bandwidth requirements will occur as true multi-media mail becomes available. Animated graphics, digitized speech and full video all require high bandwidth. Since these messages have high storage needs, only a pointer to the multi-media message may be transmitted for notification purposes. The actual contents may be sent on demand, and in real time when the reader actually refers to it.

5. Bulletin boards and asynchronous conferencing

5.1 Introduction

Bulletin boards emphasise allowing people to post messages to the on-line community , while asynchronous conferencing enables one to create, join, and actively participate in an on-going discussion of some topic of interest.

Electronic mail is used most often for one-to-one communication, and to a limited extent for one-to-many communication. Mail is sent directly to a recipient through their address. Bulletin boards, on the other hand, act as repositories for one-to-many messages, usually organized by subject matter. A person posts a message to a subject folder that is forwarded to the on-line community at large. Unlike electronic mail, the recipients are not addressed directly. Instead, anyone interested in that subject may choose to read it.

Asynchronous computer conferencing differs in emphasis from bulletin boards by promoting many-to-many communications[2]. Conferencing systems emphasise allowing people to create, join, and actively participate in on-going asynchronous discussions in topics of common interest (Baecker and Buxton, 1987). Messages and commentaries are organized by topic, with individual postings usually sorted in chronological order.

Existing systems often support both asynchronous conferencing and bulletin boards, with little differentiation between the two. Consequently, this discussion will not distinguish between the two any further . Electronic mail is normally available as well (usually as a sub-system), so community members can communicate directly as a follow-on of a posting.

Asynchronous conferencing/bulletin boards can provide considerable benefit to a research and development community. A few salient points are listed below.

1. It is a considerably faster method of distributing information than paper.
2. It fosters discussion and creativity between participants.
3. It keeps everyone well informed about current activities and topics of interest.
4. People can join and actively participate with discussions even though they may not know anyone in the group personally.
5. People may respond directly to a person’s posting via integrated electronic mail rather than to the group.
6. People may broadcast information requests over the bulletin board. For example “Does anyone know if there are any published articles on the topic of...” or “Has anyone seen this particular problem and found out how to solve it”?
7. People may post abstracts of work in progress, internal reports or publications in press, allowing interested parties to request the full paper.

---

[2]“Computer conferencing” is the popular term for asynchronous conferencing. However, it is a poor term, for the name itself does not differentiate between the ways that computers can support conferencing (eg face to face, real-time, and so on).

This section will describe the standard functions that are usually provided by asynchronous conferencing systems, raising several issues in current and future designs. Limitations and unknowns are noted, and network requirements discussed.

5.2 Functions and issues.

Current systems offer a wide range of functionality to the community. Some standard features are listed below.

Organization. Articles are organized into groups representing topics of interests or discussions.

Search facilities. People can find items of interest by searching for items in article headers (dates, people and topic names, words, topics, keyword lists), and for strings appearing within a posting.

Sorting capabilities. People can sort the message headings within a particular conference in several ways. Some examples are chronological order, alphabetic order of authors, alphabetic order of subject header.

Connections through electronic mail. People can respond directly to the author of a posting without leaving the conferencing system.

Access control. Access control may be provided so that particular groups are moderated by one person, or have a restricted membership, or control who the active participants are (ie the ones who can post to the group) separately from who is allowed to read the postings.

Information management. To remove the clutter found in large conferencing systems, people can un-subscribe to particular groups and delete unwanted postings, removing them from view. The system automatically archives those postings that are read but not deleted.

An important architectural issue with conferencing systems is where the information actually resides and how people access it. With a *centralized* architecture, postings and the group structure is maintained on one computer. People access the system through remote dial-up. With a *distributed* architecture, postings are broadcast to all subscribers. The consequences are two-fold. First, centralized architectures are often complete services that tightly couple the availability of the postings with the user interface employed to present the messages to the user. Distributed architectures, on the other hand, often have several interfaces that take advantage of the postings being distributed. Second, it is more difficult to capture the relationships of postings within a distributed architecture due to synchronization problems. For example, one site may receive a reply to a message before receiving the original message. This issue is particularly relevant when the inter-relationship between messages is important. An example is a group that wishes to separate a discussion’s central theme from side-arguments.

The main advantages of asynchronous conferencing is its ability to bring together a large number of geographically separated people with common interest over an extended period of time (Jackson, 1989). Because it is asynchronous, people can contribute to the “conversation” at their leisure. Because it is text based, old material can be reviewed as a structured record of messages (Jackson, 1989).

The asynchronous aspect also provides for possibilities of non-sequentiality within the conference. For example, consider the temporal sequence, shown in the figure below, where person A initiates some topic, B and C reply directly to A and D alters the topic. A new person E now reading this chain of messages may reply to A’s original posting, comment on C’s reply, and then discuss D’s topic change.

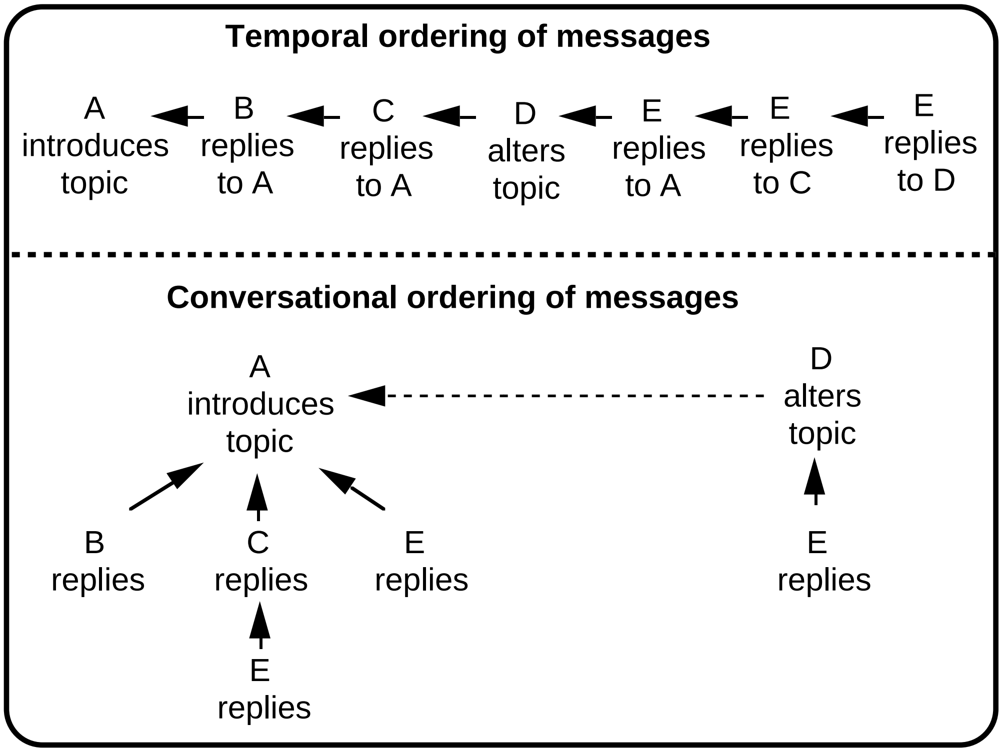

The figure also shows a logical re-construction of this simple conversation as a set of topics and responses . As can be seen, the resulting conversational pattern is quite a bit different from the temporal one.

The difference between the conversational and temporal organization of computer conferences highlights several needs of the people using the system. As Jackson writes,

> …conversations can grow and diverge into complex non-linear structures, so participants need ways to follow branching discussion paths, and to ascertain the context of any particular item in a conversation.
>
> (Jackson, 1989)

If the system structures messages by their temporal arrival (as most commercial systems do), it is increasingly difficult for a reader to:

- follow conversational paths;
- understand the context of replies;
- recognize, pursue and/or prune conversations on divergent conversations;
- get a feel for the gestalt of the conversation.

As a result, many conversations have a fragmented feel about them. Some commercially available conferencing systems pay some service to the conversational ordering of messages. For example, *Participate* is a text-based service that allows each note in a conference to become the root of another conference (Stevens, 1986). More interesting is *Banyan*, a research prototype of an asynchronous conferencing system (Jackson, 1989). Banyan represents conversations and messages as nodes within a graphical hypertext system. Links within the nodes explicitly indicate the logical ordering of the conversation. Not only can one see the context of a reply (backward chaining), but one can also inspect the original message and follow replies from that point on (forward chaining). Banyan also allows people to create views on temporary collections of messages resulting from searches or queries. People may hide uninteresting topics. Banyon is, in fact, similar to Object Lens, for they are both based upon an underlying hypertext model of message storing (Lai and Malone, 1988).

The previously discussed *Andrew Messaging System* is worth revisiting for the novel and interesting features in its bulletin board system (Borenstein and Thyberg, 1988).

- As with email, all *Andrew* messages may comprise one or more multi-media objects.
- A person is allowed to create a privately editable, publicly readable electronic “magazine”. Borenstein and Thyborg give an example of a user with a strong interest in music compiling key articles from several music-related newsgroups into a single magazine (Borenstein and Thyberg, 1988).
- Protection mechanisms permit creation of bulletin boards with a variety of access permissions. This allows bulletin boards to reflect particular needs. Examples include:
  - group oriented: readable and writable by only a designated group of people;
  - administrative and advisory: writable by all, readable by few;
  - secretary support of email: postable by a few email recipients, readable by one.
- Some messages can, when read, invoke a message-specific interaction with the user. These have been used for system-facilitated voting, and for return-receipt confirmations from the reader.

Andrew also shares several common features with Object Lens.

- Users may specify how incoming mail is to be filtered (eg deleted, re-directed, placed in special folders, and so on).
- Message templates are available for replies.

The widespread use and the extremely large content size of even primitive asynchronous conferencing systems indicates how important they are. Unlike all other forms of computer-supported cooperative work mentioned so for, conferencing systems travel well beyond the small group by bringing together an extended community of people with common interests.

Semi-formal messaging, described in the last section, can be extended as semi-formal conferencing. The best example is gIBIS, a system that captures early design deliberations on large complex problems (Conklin, 1988). It is based upon the Issue Based Information Design (IBIS) methodology that views design as a rhetorical process, with a set of issues that can be generalized, specialized, responded to, questioned, argued and so on. As with Object Lens, gIBIS is based to a large part on semi-structured messages. Through its well-designed interface, participants propose and respond to issues in structured ways that eliminate unconstructive moves such as name-calling and argument by repetition. A key aspect of gIBIS is that all messages are captured and structured within a group-accessible hypertext database. As with most hypertext systems, the relationship between text fragments—the particular arguments made—are illustrated and manipulated graphically by the user.

5.3 Limitations and unknowns

There are several issues that plague conferencing systems. The first is the user interface. A conferencing system will only work if many people can access it. Because of the wide variety of display hardware and the lack of standards, this has restricted most systems to glass-teletype ascii-based interfaces, a poor medium for representing the rich inter-connections between postings. Fortunately, the advent of X window standard will likely see new conferencing interfaces window based upon graphical systems.

The second issue is the difficulty of handling and filtering the tremendous amount of postings that become available in some conferencing system. In the Unix-based Usenet, for example, not only are hundreds of news groups available, but each group may have tens to hundreds of postings made to it daily or weekly. To make matters worse, un-moderated groups are often subject to “junk” postings (often called flames) that add little to the group’s discussion. New subscribers are often over-whelmed, while even faithful readers are often hard-pressed to keep up to date. The result is that most users subscribe to just a few selected groups. When they fall behind in their reading, they often prefer to skip over old material so that they can keep up with the new. Relevant postings may be lost as a consequence. The tentative solution is to incorporate intelligent filtering into the system, so that un-interesting articles are pruned and interesting ones made prominent. While simple filters exist in some systems, intelligent filters must wait until ideas similar to those in Object Lens (Lai and Malone, 1988) are commercialized.

The third issue arises from semi-formal conferencing. Conklin and Begeman evaluated and criticized the gIBIS system work upon preliminary observations of its use (Conklin, 1987). On the positive side, users found the IBIS formalism (and the gIBIS tool) to be a powerful method for research thinking and design deliberation, especially for detecting incompleteness and inconsistency in thinking and rhetoric that added little to the argumentation. Users also found the semi-structured aspects of gIBIS aided message composition, search, and comprehension.

Several negative aspects arose with the IBIS formalism as implemented by gIBIS (Conklin, 1987). IBIS sometimes proved inadequate for capturing all aspects of the problem being considered, and it did not deal well with the vague, contradictory and incomplete thoughts usually expressed in the early stages of a problem. Meta-level communications is another issue. While gIBIS supports the substantive part of the work, it does not support annotative comments about the work, nor procedural comments about the communications itself. In spite of these criticisms, we should remember that gIBIS is an early prototype and has much room for improvement.

5.4 Network requirements

The network requirements for conferencing are range from low to extreme. As with electronic mail, low-bandwidth systems can handle even large numbers of postings of limited size as long as immediate delivery is not necessary.

Multi-media postings require significant bandwidth. Consider a centralized architecture, where the user must log on to the remotely-located system. Since postings are transmitted only when they are selected by the user from the subject heading, they must be sent in real-time. The network must be able to transmit the posting with little delay, for excessive waiting would deter the reader from scanning all but the most interesting articles. With a distributed architecture, articles must be broadcast to all subscribed sites as they are posted.

Given a large number of sites and a high volume of postings, considerable demands may be made of the network.

6. Access, Monitoring and Operation of Scarce and Distributed Equipment

6.1 Introduction

Real-time monitoring and operation of distributed equipment is now commonplace within industrial settings, usually through real-time systems. Systems include supervisory control and data acquisition, tele-operation, and tele-robotics. Availability remote computational power is discussed here as well. All are highly relevant to researchers and developers who require access to remote equipment.

1. Equipment is a shared resource available to non-collaborators. This often occurs when high equipment costs must be amortized over a large number of people (for example, a super-computer or a radio telescope) and when the equipment is rare (perhaps due to it being a novel technology).
2. Intensive computation may be distributed in parallel to idle computers in the network.
3. Some instrumentation is inaccessible due to its location (underground, space, deep sea) and hazardous environments (nuclear power plants).
4. Members of a distributed team monitor and control equipment located at one collaborator's site. One example is simultaneous engineering.
5. A scientific community observes an experiment in progress (one example recently occurred when astronomers from around the world gained limited access to the Voyageur II's Uranus data as it was transmitted to Earth). This is valuable to obtain early community reaction to methodology and analysis, or when the data is so complex that good analysis requires study of data by multiple independent experts.
6. Remote observation and operation is an extension of "tele-data" by including observation and control of physical devices as well as software. Not only can people see what the other is referring too, but have greater opportunity to interact through these devices.
7. Hands-on distance education and support becomes possible. This is especially important for keeping the community up to date on the latest technology, and for the on-going support requirements that occur when one consults an expert.
8. Maintenance of specialized equipment may occur through specialists remotely monitoring and servicing the system. This is important when the users of the equipment are not expected to have the expertise to maintain it.
9. Programmable robotic systems allow flexible re-tooling and re-configuration of distant systems so that it can handle a variety of situations.

6.2 Computational Access

People often need to access and work on remote computers. They may be working through a terminal with no compute power, their own machine may be too slow, or necessary hardware and information is only available on the remote system. This sub-section reviews the several ways that people access remote computers.

ASCII-based interfaces. The simplest way to access a remote computer is through a modem or network. Since most personal computers and workstations can emulate ASCII terminals, the person simply logs in as normal once the connection has been made. Typical connection speeds through modems are 1200 to 2400 baud, with modem speeds up to 9600 baud now becoming available. The later is fairly close to the line speeds of terminals hooked directly into the computer. Networked workstations can, of course, offer greater speeds for remote logins.

Networked window systems. ASCII-based systems are reaching the end of their lifetime. Window systems running graphical applications are now becoming the standard platform for human/computer interfaces. Until quite recently, these window systems required a person to be logged in directly to a bit-mapped workstation; they would not run over a network. However modern networked-based window systems, such as Sun Microsystem's NeWS and the Open Software Foundations X Windows, have removed this limitation. A person may have several windows on display, with each perhaps running an application being executed on a completely different computer. In essence, the communication protocol for the user interface has shifted away from low-bandwidth ASCII to a higher-bandwidth graphics standard. Consequently, the connect speeds once adequate for ASCII terminal emulation are painfully slow for emulating window systems.

Distributed computation. While workstations are becoming increasingly powerful, the demands being made of them are also rising in complexity. Distributed computing is one method of overcoming this problem. At its simplest, a user connects to a powerful computer (say a supercomputer), using a less-powerful computer (such as a workstation) to handle the more mundane computational tasks. But increasing attention is now being paid to *parallel computation*, where many computers are linked together to solve a problem that has been functionally decomposed into portions that can be executed in parallel. A low-cost approach to parallelism is to have one's machine search for idle or under-utilized computers on the network and to distribute the decomposed problem accordingly. Another approach is to use the dedicated parallel computers now coming onto the market. In spite of these advances, the bottleneck to parallel computation is the network speed. The smaller the parallel units to be distributed, the greater the communication demands.

6.3 Technologies for monitoring and operating remote equipment

There are a variety of technologies relevant to monitoring and controlling remote equipment. Many are now in use in specialized settings (ie heavy industry, research laboratories, space), but their use should become more available to the community as costs and availability decrease.

At the heart of control and operation of remote equipment is the area of *real-time systems*. As its name indicates, these systems demand response in real time. Significant time delays for relaying information to and from a controller and the equipment is not acceptable. Several different flavours of real-time systems exists.

SCADA—Supervisory Control and Data Acquisition systems are the slow end of real-time systems. They commonly operate on low speed networks (such as telephone lines with modems). SCADA systems are responsible for monitoring and occasionally over-riding the (usually) automatic operation of remote sites—the latter is called supervisory control. In typical use, *remote terminal units* send data points from tens or thousands of widely-distributed and heavily instrumented sites to a centralized database, where they are checked against pre-defined constraints. Alarm conditions are raised on an operator's console when allowable bounds are passed. Operators then view the collected data points through mimic diagrams displayed on a screen, apply their knowledge to control the remote system from their consoles, and see the resulting changes almost immediately. Examples of SCADA system use range from operation of elevators within a building, city-wide electrical utilities, and international gas pipeline system control.

As applications become less tolerant of time delays and the need for control becomes more rigourous, we move closer to the heart of real-time systems where SCADA techniques are no longer acceptable. We will differentiate “true” real-time systems from SCADA by its need for continual rather than supervisory control. Time frames also differ. Network transmission of seconds or even minutes is often acceptable for SCADA; real-time systems often require tolerance levels of milli-seconds. Real-time systems are used for a variety of purposes, from close monitoring and control of critical operations within a plant (such as a nuclear power plant) to tele-operations of robots.

Master-slave tele-operation covers any manipulation system that is remotely controlled by a human operator in response to sensory information transmitted from the workplace (Scott, 1984). One example is direct human control of a robot’s actions in real-time, usually by moving a device that mimics the motion of the actual robot. These systems are required when the need for human involvement is high, the working environment is remotely located or hazardous, and when the size/strength/manipulation differential between human and machine is great. Because the robot is not autonomous, a high degree of operator feedback is needed.

Tele-robotics covers computer-controlled tele-operation of robots, and is one of the most promising application areas to support distributed research and development. Because robots are generic in operation (such as robot arms), they can be programmed to perform a variety of specific tasks related to specific needs. Tele-robotics are already being seen in *flexible manufacturing*, a robot-based system for the manufacture of diverse products.

There are two basic methods of tele-robotic operation.

1. A computer sends high-level instructions or downloads programs to a remote robot, so that many low-level actions are handled by a controller local to the robot. In this situation, the robot is considered semi-autonomous. It can react to immediate situations (such as dealing with obstacles) without referring back to the distant computer.
2. The centralized computer directly controls the remote robot. This may be necessary when there is not enough compute power available to the robot, or when part of the computer’s task is to synchronize the robot’s actions with other machines it may be controlling (such as other robots or an assembly line).

The semi-autonomous robot is the goal of much robotics research, for the greater the on-board intelligence of the robot, the less bandwidth and rapid reaction time required by the distant controller. This is especially critical in space robotics, where speed of light constraints introduce unacceptable response for distant control. Even so, semi-autonomous systems are usually linked to a central system so that humans can monitor and intervene in the operation, and for data collection purposes.

6.4 Network requirements

Computational access requires modest network speeds. Normal telephone speeds are acceptable for ASCII-terminal emulation; and a 1.5Mbps network is adequate for window-based interfaces. Network requirements for parallel systems can range from low to high, depending upon the size of the computational unit and the amount of communication between these units.

Shifting the emphasis towards monitoring and operation, existing SCADA systems require relatively low-bandwidths. They are well served by current telephone lines and would certainly be well covered by the 1.5Mbps Network. True real-time systems, especially those controlling robots, have much stricter bandwidth demands and time tolerances. Speed is critical. In master-server tele-operations, the operator may need a high-resolution video image of the remote robot to see what is going on. Additionally, the feedback cycle between action, consequence, and notification cycle must be short: the maximum tolerable feedback delay for effective real-time control of the robot by the human is one tenth of a second (Scott, 1984).

Computer controlled tele-robotics usually require high-bandwidth reflecting the data rates for control and feedback of servers, and perhaps information transmission by any attached vision systems. The preliminary NASA/NBS Standard Reference Model for Tele-robotic Control System Architecture (NASREM) describes timing tolerances of a few milli-seconds for synchronizing activity, especially at the servo level (Albus, McCain and Lumia, 1989). Even semi-autonomous robots may require bandwidth for monitoring (perhaps through video), data collection, transmission of high-level control information, and synchronization with other devices. Other criteria are, of course, system dependant.

Robotics have a good migration path to a nation-wide network, simply because they already rely on software supervision and often use high-speed local area networks for communication.

7. Digital Libraries

7.1 Introduction

Discoveries in research, technology, and development are exploding. Not only is the amount of information produced growing, but it is becoming increasingly more specialized. As far back as 1945, Vannevar Bush described the dilemma we are now surmounted with:

> …there is increased evidence that we are being bogged down today as specialization extends. The investigator is staggered by the findings and conclusions of thousands of other workers—conclusions which he cannot find time to grasp, much less to remember, as they appear. Yet specialization becomes increasingly necessary for progress, and the effort to bridge between disciplines is correspondingly superficial.
>
> (Bush, 1945)

How will information be handled in the next decade? It is clear that traditional technologies, such as paper books and journals, are unable to accommodate well a researcher's need for timely and relevant information. Review and publication processes for conference proceedings and journals normally span the better part of a year, while book production is even longer. Information searching and retrieval is difficult even with one's area of specialty. As an expert moves outside of his specialty, it becomes almost impossible to find, let alone be informed of the key related material.

We require new ways of manipulating information. At the heart of these methods will be the *digital library*—a computational way of storing, relating and retrieving information. In many respects, digital libraries are related to bulletin boards/asynchronous conferencing system in that both supply similar benefits. They are two ends of a spectrum, with bulletin boards addressing the *process* of creating information (drafts, discussions), and digital libraries addressing the *archiving*, *relating* and *access* of information. A digital library has potential to provide many benefits to research and development.

1. Researchers can see their works published immediately.
2. Published works become available to the community immediately.
3. People may discover and access "grey" publications—technical reports published internally in an institution—through the library. Grey reports are currently difficult to find out about, and even more difficult to acquire.
4. People may review information presented in alternate and more relevant forms. The technology of publication is extended to include multi-media documents, such as video, sound, animation, active programs, and so on.
5. A specialist can search, filter and retrieve relevant information from large databases. Search and retrieval facilities need not be restricted to specialized archival sources (as they are now).
6. Since there is no notion of being "out of print," researchers may obtain historical articles as easily as contemporary ones.
7. Readers can follow cross-references easily (such as bibliographic citations) by retrieving the reference directly.

8. Similarly, readers may review attached commentaries, rebuttals or extensions to articles.
9. A researcher may personalize the information by adding notes and links to relevant material. Additionally, one may gain a new perspective to the information by viewing another person's notes and links.

The list can grow almost indefinitely. The impact of a digital library to society is potentially as beneficial as the introduction of the printing press.

This section will briefly describe the conventional approaches related to digital libraries, and then discuss the hypertext approach. Limitations and unknowns are raised, and the network requirements provided.

7.2 Conventional approaches.

There are two standard ways of storing information on computers—as hierarchical file systems and within a highly-structured database. As both methods are likely well-understood by the reader of this document, they will not be described. Rather, their use, suitability and influence to the digital library will be discussed through their roles as a digital repository and a bibliographic database.

Digital repositories are file systems maintained by an organization such that files can be accessed and copied by people at another organization. They are currently used by institutions that make information—usually software and documentation—available in the public domain. Access is normally through a network such as the Internet. Consider the *Free Software Foundation*, an organization that maintains a digital repository of Unix-based software. Although people can buy a tape of the library at a nominal cost, one may use the file transfer protocol (ftp) across the Internet to list and transfer part or all of the files from the foundation's machine to one's local machine. Another digital repository is available through the White Sands Missile Range, who maintain ADA-related software and documents. Yet another is a software repository held by Sun Microsystems, who allow people to request and transfer software through an automated electronic mail server. Many universities also have similar repositories to encourage information exchange and cross-fertilization—in essence, a distributed library system.

Of course, digital repositories are extremely limited. They do not have any search capability, and people only find out about a repository's contents by mundane means, such as an announcement on a bulletin board or by remotely accessing and searching around the file hierarchy. Still, they provide a simple mechanism for distributing information at relatively low cost.

Bibliographic databases are databases containing citations of publications. Most are limited in scope. They either reflect a special topic of interest or the holdings of a physical library. Although abstracts of publications are sometimes included in a citation, the full text is rarely available.

A person would use a bibliographic database for several reasons. First, one may be interested in *searching* for all titles related to a particular, quite specific topic. If abstracts are available, one could then scan them for direct relevancy. Second, one may wish to *locate* a publication at a physical library, possibly arranging to receive a hardcopy through an inter-library loan.

As a highly specialized form of a digital library, bibliographic databases have several deficiencies. First, the skills required at choosing appropriate search terms usually requires a librarian to act as an intermediary. Second, they rarely provide the actual publication, and they do not list relevant cross-references and commentaries. Third, grey publications are rarely included. Fourth, they are centralized. Users must usually log onto them remotely over a network or modem via a terminal emulator, with little direct capability for transferring the found information to one's local machine.

7.3 Hypertext and Hypermedia

Neither files systems nor traditional databases are an adequate backbone for future digital libraries. While files lack the structure needed, databases are too rigid in the structure they provide. An alternative lies somewhere between these extremes—hypertext and hypermedia.

Hypertext extends the notion of a document beyond sequential text by allowing complex interwoven structures to be created and manipulated by linking text fragments. The fundamental idea is simple: links can be added anywhere in the text database which, when followed, will transport the reader to another location. Associating types with links extends the power for enhancing semi-structured access to a document's contents, with instant availability of:

- related information;
- chasing references;
- visualization of the information structure;
- examples illustrating the text;
- rich searching and indexing facilities;
- selective and personal in-depth explorations;
- annotations comprising definitions, footnotes and asides;
- marked trails of the path through the information taken by the reader;
- guided annotated tours defining a "tour leader's" perspective and comments (Trigg, 1988);
- convenient opportunities for activities such as adding personal annotations and place-marking.

Hypertext becomes *hypermedia* when any media form can be used and linked into the document. An author's point may be annotated with an instantly accessible image, sound track, or video clip. Sometimes active sections (ie a running program) may be incorporated into otherwise passive documents to permit user interaction. When this ability is added, hypermedia becomes a rich new metaphor for interacting with computers and file stores.

Fundamental to hypertext and hypermedia is the user interface used to navigate information and traverse links. Most interfaces are based upon *browsers*, where node contents are displayed in a window and links to other nodes are marked through in-text icons. Browsers usually include some type of overview facility that provides a map through the information, usually by displaying the network graphically. Readers typically search for information by following the links embedded in the text, by searching the network, and by selecting an item directly from the overview map (Conklin, 1987).

There are now a large number of research and commercial hypertext systems. Early systems were envisioned as macro literary systems that focus on the integration of colossal volumes of information (Conklin, 1987). A few are listed below.

1. The first is Bush's un-implemented vision of the *Memex* system (Bush, 1945). His idea was not based upon digital technology; rather, large stores of scientific literature would be linked through microfilm and photocells.
2. In the 1960's and 1970's, Doug Engelbart proposed and built NLS/Augment, a system that emphasises an environment for "knowledge workers" (Engelbart and English, 1968). Not only was NLS/Augment important as an early implementation, but it still stands out as an environment that integrates hypertext, communication, document processing, idea exchange and information management.
3. Perhaps the most well-known macro-literary system is Project Xanadu, developed and now being implemented by Ted Nelson (Nelson, 1987). Nelson's goal is ambitious, for he sees placing the entire world's literary corpus on line (Conklin, 1987). Rather than concentrate on the user interface, Xanadu investigates the database requirements needed to maintain such a large, heavily linked information store, as well as providing mechanisms for accounting and royalty distribution.

While the early visions of macro-literary systems are suitable to a large, centralized digital library, most current commercial systems are not, for they are designed to accommodate small information structures equivalent to (say) a conventional book. But hypertext will continue to grow over the next decade, and will catch up to the early visions of Bush, Engelbart and Nelson.

7.4 Limitations and unknowns

As with any new form of literary presentation, there are many problems yet to be solved in hypertext and hypermedia. While some are due to current shortcomings of hypertext interfaces, two critical issues arrive from human behaviour: disorientation and cognitive overhead (Conklin, 1987).

Disorientation is the "lost in hyperspace" problem that occurs when people lose track of where they are in the network, or are unclear of how to get to some other place in the network. The difficulty arises because, unlike sequential text, readers are free to wander off in any direction. Technical solutions may overcome some difficulties. Graphical browsers may trace and visually show a person's location in the hypertext document, perhaps distinguishing between mainline progress and side-excursions. Sophisticated query facilities similar to those available in databases can help searchers find and navigate to desired nodes. Filters may remove unnecessary clutter from display, showing only links of interest to the reader. "Meta-structure" may be added to the hypertext system as "guided tours" through selected paths (Trigg, 1988).

Cognitive overhead arises from the additional effort and concentration necessary to decide if links are worth following, and when one maintains several tasks or "trails" through a document at one time (Conklin, 1987). Is a side path worth taking? How does one switch conceptual context back to the main path after chasing one or more side paths? Again, technical solution may ease the problem. Brief explanations of node contents may be made available from the link itself to help a reader decide if a link is worth pursuing. Graphical browsers can help people track excursions, helping them maintain and return to previous contexts.

7.5 Network requirements

Conventional bibliographic databases have low network requirements due to the small amount of information generated; they now run adequately over modems. In marked contrast, file transfer from digital repositories can involve many megabytes of information. Some repositories have software releases available of 50 megabytes or more! Fortunately file transfer is rarely time critical, and are still well-served by low-bandwidth lines.

In contrast, hypertext can require significantly more bandwidth. If a hypertext system is maintained at a centralized location, then users must access it through the network. Information retrieval is time-critical. If a reader chases a link, the node contents should appear as quickly as possible. Delays will deter readers from chasing paths, contradicting the intention of hypertext. While moderate speed lines may suffice for pure text systems, hypermedia nodes would require the bandwidth needed for rapid transmission of audio, graphics, video, and so on. Additional bandwidth is also needed to drive the graphical user interface employed by most hypertext systems.

8. Application areas

8.1 Introduction

This report began by separating the primary research and collaboration functions from the collaboration tools required to implement those functions. The previous sections have concentrated on these collaboration tools, mentioning several of the functions as examples that can benefit a distributed research and development community. This section will highlight three application areas to indicate several ways that the collaboration tools can be applied: distance education, group decision support systems, and joint authoring. It closes with a note describing the need for the network to bootstrap collaborations.

8.2 Distance education

Certainly one of the most critical issues facing a research and development community is in training and continuing education. Excellence of work demands that researchers continue to gain familiarity with new equipment, learn new trends in their own area of expertise, and to extend their horizon towards other disciplines. Many of the systems already described in this report can provide the backbone of a distance training and education system running over a network.

Current computer-aided instruction (CAI) usually revolves around learner interaction with information in a database. Courses are prepared instructional modules where learners can go through the material at their own discretion and pace. Both distributed libraries and hypermedia support CAI. Packaged courses can be distributed to remote learners through a digital library. Using some of the ideas available in hypermedia, directed courses can be augmented by the relevant scientific readings, allowing learners to pursue a topic as lightly or deeply as desired.

However, good education demands learner interaction with the instructor and other learners. By augmenting CAI with electronic mail and bulletin boards, learners may contact the course organizer for questions and clarification, or even each other for on-going “class” discussions of the material. When courseware is built as a simulation system such as SharedArk and cyberspace, the richness of the visualization and interaction and the possibility of real-time multi-student collaboration extend beyond anything now available.

Equally as exciting is the ability to hold real-time distributed lectures and seminars, where the lecture halls are the sites connected through real-time tele-conferencing. This is particularly advantageous to organizations that cannot afford to bring in prestigious speakers. By linking in to larger organizations or by cost-sharing, researchers at even small institutions can gain access to presentations made by prominent figures. Multi-way interactions between the presenter and audience, and the audience members with each other create a rich environment encouraging questions, comments and discussions. The tools of the classroom are provided through tele-data—the whiteboard, the shared notepad. Even real-time “hands-on” training could be available through tele-operation and tele-robotics.

Distance education must satisfy many criteria to work. The educational resources and technologies must be physically accessible; the material must be topical to the learner;

people involved should be collaborative-accessible (Quigley, 1989). The proposed Network is critical for achieving the first and third points.

8.3 Group decision support systems

Group decision support systems (GDSS) combines computer-supported communication and decision support technologies to facilitate formulation and solution of unstructured problems by a group of people (DeSanctis and Gallupe, 1987). The goals of GDSS are:

> …to reduce the "process loss" associated with disorganized activity, member domination, social pressure, inhibition of expression, and other difficulties commonly encountered in groups and, at the same time, to increase the efficiency and quality of the resulting group decision.
>
> (Watson, DeSanctis and Poole, 1988)

DeSanctis offers three levels of GDSS systems (DeSanctis and Gallupe, 1987).

Level 1: Remove common communication barriers to facilitate information exchange. These are the areas usually handled by the generic collaboration tools mentioned in the previous sections: tele-conferencing, email, shared workspaces, and so on.

Level 2: Provide decision modeling and group decision techniques aimed at reducing uncertainty and noise that occur in the group's decision process. These include planning tools, agenda setting, decision modeling methods (such as decision trees, risk analysis, and forecasting methods).

Level 3: Structure the group communication pattern (the rules of discourse) to one appropriate to achieving the goals of the meeting. Structured group methods include the Nominal Group and Delphi techniques, as well as the IBIS and speech act approach described earlier.

A GDSS clearly has an important role in distributed research and development, especially when inefficiencies and uncertainties related to the normal decision process are exaggerated due to distance. The collaboration tools mentioned in previous sections only fullfill the first level of GDSS; they must be extended to include the specialized techniques for structured decision analysis, and at directing and structuring the discussion. Unlike most generic communication systems, GDSS systems intervene in the group's natural decision process to help them come to a higher quality decision.

8.4 Joint authoring

As with GDSS, distributed joint authoring is a specialization and extension of the previously defined collaboration tools. Joint authoring can take many flavours. At its simplest, document drafts and commentary may be exchanged by electronic mail, or perhaps worked on in real time through tele-data. More sophisticated authoring tools would understand and support directly the nature of the interaction between authors, reviewers, and readers.

Consider Quilt, a tool for collaborative document production (Leland et al., 1988). Unlike other collaborative document systems which support only direct authoring aspects, Quilt emphasizes and supports the communication vital to good collaboration. For example, structured hypermedia links allow people to attach text and voice annotations to the document, specialized as revision suggestions, public comments, and directed messages. The necessary coordination between collaborators is enhanced via activity logging, notification and triggering mechanisms. Access permissions can be set by the author to reflect the varying roles of collaborators (as writers, commentors, reviewers), while user-customizable definitions for such things as document and annotation types make the system both flexible and extensible.

Specialized joint authoring systems can make use of many of the technologies mentioned in previous sections: real-time and asynchronous communication; bulletin boards and asynchronous conferencing for communal reaction; digital libraries acting as a centralized store of the evolving paper that is accessible by many.

8.5 Bootstrapping collaborations

Unlike applications intended for a single users or sites, collaboration tools operating over distance have a bootstrapping problem. Researchers and developers will not use systems until they are available and ubiquitous, while vendors will not build the applications until these people will commit to buying them. In essence, a critical mass of people and equipment is required to bootstrap the process.

High speed networks are currently too expensive and too difficult to obtain for the average researcher. Their use is restrained to organizations with exceptional needs. As long as the status quo remains, collaboration tools are unlikely to make any impact.

The delivery of the Network can overcome this problem in several ways. First, it will act as a testbed for research and development for the network itself, and of future collaboration tools and information technologies. Second, because of its availability to a diverse set of people, members of the community with real needs can be used in pilot studies for further testing and tuning of the collaboration applications. As the applications become closer to commercialization, their introduction to the community becomes a matter of extending the audience already addressed over the network, instead of starting afresh. The act of having the application available before-hand will allow part of the community, perhaps a critical mass, to experience, buy into, and advertise the application.

9. The Coherent Workspace

The networked nation is more than a means of overcoming distance barriers, for it has the ability to augment researchers' and developers' knowledge and activities beyond those now found traditionally. Each tool and function mentioned in this report can, by itself, provide added value. Communication vehicles will fit tasks, digital libraries will become knowledge bases; tele-robotics will extend our physical limitations. Yet the real power will arrive when tools and people are integrated together into the *Coherent Workspace*.

The Coherent Workspace occurs when people and their machines have a common knowledge and sense of purpose, with individual and group activities being well-coordinated. Coherent systems have their roots in several early visions. Vanevar Bush saw *Memex* as a way to tie information and people together through a dynamic richly-connected knowledge base (Bush, 1945). Doug Engelbart extended the Memex theme through the *Community Handbook*, a coordinated handling of a very large, complex and continually evolving knowledge base that comprehensively represents the current state of a collaborative endeavour (Engelbart and Lehtman, 1988). These include:

- *formative information* for a collaboration such as purpose, working hypotheses, goals and expectations;
- *project management* such as commitment tracking, status schedules, project and meeting management tools, and standards employed;
- *project content*, such as observations, data, and results;
- *training and education*, such as how-to methods;
- *support and incorporation of communications* to support and track the on-going dialogs about the project and contents of the database.

A coherent workspace includes these notions through four principles: persistence, inter-operability, integration, and coordination.

The *principle of persistence* covers the integrity of an activity across time and distance, across group size, and across changing group membership. Asynchronous bulletin boards are a simple but good example. The topic and on-going discussions of a "meeting" are captured, so that old and new group members can review what has happened and how the thread of the dialog has altered over time. The discussion persists even when discussants drop out, new ones come in, or existing ones change roles from, say, observers to active speakers. But persistent systems could do much more. First, they could capture all of a group's activity, independent of time and place. Second, the system could explicitly maintain the goals of the group, perhaps moderating the activity to keep it on track.

Consider an anecdote of a group of engineers designing brakes for an automobile[3]. The design and final part produced included a burr of metal attached to the assembly to help dissipate heat build-up. Several years later, a new team re-designed the brake. Although they had the final blueprints of the earlier brake, they did not have the reasons for the design decisions made. They removed the seemingly non-functional burr for cosmetic reasons. The result was a recall of all cars with the new, now deficient brake. A persistent system, on the other hand, would have captured the on-going discussion, the decisions

[3]Example provided by Dr. Marilyn Mantei, University of Toronto.

made, and the reasons behind the decisions, and extended those across the "null meeting" period until the new group met.

The *principle of inter-operability* applies when independent systems operate well together. This implies standards for communication, common representations of problems, standard interfaces, machine independence, and so on. Commercial applications are now beginning to follow inter-operability, at least on a low level. In the user interface domain, look and feel standards are now available, such as Unix International's *Open Look* and the Open Software Foundation's *Motif*. Machine independence is provided by networked-based window and graphics protocols such as X windows, which allows a program running on one vendor's workstation to be viewed on another vendor's machine. Inter-program communication is also making some headway. For example, many Apple Macintosh programs now follow an interchange format allowing items copied from one program's view to be pasted into another program's view without either being aware of the other. In artificial intelligence, blackboard systems allow programs to make use of and to update a common knowledge representation.

Inter-operability becomes the *principle of integration* when sub-systems are tightly coupled. Updating one component would directly and immediately be reflected in the update of another. Perhaps all components would be built upon a single underlying representation of the on-going activity, with the components merely being different views and perspectives of the knowledge base and group.

Finally, the *principle of coordination* integrates all activity in the sense of group awareness, common knowledge, purpose, relevant training, and coordination of activities. It maintains the goals of the group, schedules tasks, commitments and dependencies, the progress to date. It keeps members up to date on the relevant progress of the group and of the external community.

Although elusive, the coherent workspace is under active development. One example is a program now being developed at McDonnell Douglas, a complete information system that supports very large aerospace programs (thousands of people), over the whole life-cycle of the project (e.g., three decades) of extremely complex products[4].

[4]Example provided by Doug Engelbart.

10. References

Albus, J. S., McCain, H. G. and Lumia, R. (1989) “NASA/NBS Standard reference model for telerobot control system architecture (NASREM).” NIST Technical Note 1235, Robot Systems Division, Center for Manufacturing Engineering, National Institute of Standards and Technology, Gaithersburg, MD,

Allen, T. J. (1977) *Managing the flow of technology.* Cambridge, Mass, MIT Press,

Baecker, R. M. and Buxton, W. A. S. (1987) *Readings in human-computer interaction: A multidisciplinary approach.* Los Altos, California, Morgan Kaufmann,

Bair, J. H. and Gale, S. (1988) “An investigation of the Coordinator as an example of computer supported cooperative work.” Hewlett Packard Laboratories, California, Unpublished.

Borenstein, N., Everhart, C., Rosenberg, J. and Stoller, A. (1988) “A multi-media message system for Andrew.” In *Proceedings of the USENIX Winter Conference,* 37-42, Dallas, Texas, February 9-12,

Borenstein, N. and Thyberg, A. (1988) “Cooperative work in the Andrew message system.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 306-323, Portland, Oregon, September 26-28, ACM Press.

Bush, V. (1945) “As we may think.” *Atlantic Monthly*, **176**(1), 101-108, June. Reprinted in Greif, 1988.

Chang, E. (1986) “Participant systems.” *Future Computing Systems*, **1**(3), 253-270,

Chang, E. (1989) “Protocols for group coordination in participant systems.” In *The Structure of Multimodal Dialogue*, edition). 229-247, North-Holland, Taylor, M. M., Neel, F. and Bouwhuis, D. G., Elsevier Science Publishers B.V.

Chapanis, A. (1975) “Interactive human communication.” *Scientific American*, **232**(3), 36-42, Reprinted in Greif, 1988.

Clark, W. J. (1989) “The MIAC audiographic conferencing system: A practical implementation of the audiovisual service infrastructure.” In *Multimedia '89: Proceedings of the 2nd International Workshop on Multimedia Communications*, Montebello, Quebec, April 20-23,

Conklin, J. (1987) “Hypertext: An introduction and survey.” *IEEE Computer*, **20**(9), 17-41, Reprinted in Greif, 1988.

Conklin, J. (1988) “gIBIS: A hypertext tool for exploratory policy discussion.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 140-152, Portland, Oregon, September 26-28, ACM Press.

Corey, D., Abel, M., Bulick, S. and Schmidt, J. (1989) “Multi-media communication: The US WEST advanced technologies prototype system.” In *Fifth IEEE Workshop on Telematics*, Denver, Colorado, September 17-21, Submitted paper.

DeSanctis, G. and Gallupe, R. B. (1987) “A foundation for the study of group decision support systems.” *Management Science*, **33**(5), 589-609, May.

Egido, C. (1988) “Video conferencing as a technology to support group work: A review of its failures.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 13-24, Portland, Oregon, September 26-28, ACM Press.

Engelbart, D. and English, W. (1968) “A research center for augmenting human intellect.” In *Proceedings of the Fall Joint Computing Conference,* edition). **33**, 395-410, Montvale, NY, Fall, AFIPS Press. Reprinted in Greif, 1988.

Engelbart, D. and Lehtman, H. (1988) “Working together.” *Byte Magazine*, , December.

Ensor, J. R. (1989) “Rapport: A multimedia conferencing system.” *The ACM SIGGRAPH Video Review Supplement to Computer Graphics,* **45**(5), ACM Press, Baltimore, MD. Videotape.

Ensor, J. R., Ahuja, S. R., Horn, D. N. and Lucco, S. E. (1988) “The Rapport multimedia conferencing system--a software overview.” In *Proceedings of the 2nd IEEE Conference of Computer Workstations*, 52-58, Santa Clara, March 7-10,

Eveland, J. D. and Bikson, T. (1986) “Evolving electronic communication networks: An empirical assessment.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 91-101, Austin, Texas, December 3-5, ACM Press.

Eveland, J. D. and Bikson, T. K. (1988) “Work group structures and computer support: A field experment.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 324-343, Portland, Oregon, September 26-28, ACM Press.

Farallon (1988) “Timbuktu user's guide.” edition). Farallon Computing Inc., Berkely, California,

Fish, R. S. (1989) “Cruiser: A multi-media system for social browsing.” *The ACM SIGGRAPH Video Review Supplement to Computer Graphics*, **45**(6), ACM Press, Baltimore, MD. Videotape.

Foster, G. and Stefik, M. (1986) “Cognoter: Theory and practice of a Colab-orative tool.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 7-15, Austin, Texas, December 3-5, ACM Press.

Goodman, G. and Abel, M. (1987) “Communication and collaboration: Facilitating cooperative work through communciation.” *Office: Technology and People*, **3**(2), 129-146,

Gould, J. D. and Boies, S. J. (1983) “Human factors challenges in creating a principal support office system — The speech filing system approach.” *ACM Transactions on Offcie Information Systems*, **1**(4), 273—298, Also published in Baecker & Buxton, 1987.

Gould, J. D. and Boies, S. J. (1984) “Speech filing — An office system for principles.” *IBM Systems Journal*, **23**(1), 65-81, Also published in Baecker & Buxton, 1987.

Greenberg, S. (1990) “Sharing views and interactions with single-user applications.” In *Proceedings of the Conference on Office Information Systems*, Boston, April, Also available as Research Report, Advanced Technologies, Alberta Research Council, Calgary, Alberta, Canada.

Greenberg, S. and Chang, E. (1989) “Computer support for real time collaborative work.” In *Proceedings of the Conference on Numerical Mathematics and Computing*, Winnipeg, Manitoba, September 28-30, Available in Congressus Numerantium vol 74 and 75.

Greif, I. (1988) *Computer-supported cooperative work: A book of readings.* San Mateo, California, Morgan Kaufmann Publishers Inc,

Gust, P. (1989) “Multi-user interfaces for extended group collaboration.” In *Proceedings of the Groupware Technology Workshop*, Xerox PARC, Palo Alto, California, August,

Guttman, S. (1989) “Farallon: Building the ultimate network.” *MicroTimes*, **6**(7), July.

Harrison, S. (1989) “The office design project.” *The ACM SIGGRAPH Video Review Supplement to Computer Graphics*, **45**(4), Videotape.

Jackson, S. L. (1989) “Hypertext for computer conferencing.” Research Report STL-89-8, Hewlett-Packard Laboratories, Palo Alto, Califiornia, June.

Johansen, R. and Bullen, C. (1984) “Thinking ahead: What to expect from teleconferencing.” *Harvard Business Review*, , 4-10, March/April. Reprinted in Greif, 1988.

Kraut, R., Egido, C. and Galegher, J. (1988a) “Patterns of contact and communication in scientific collaboration.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 1-12, Portland, Oregon, September 26-28, ACM Press.

Kraut, R., Galegher, J. and Egido, C. (1988b) “Relationships and tasks in scientific research collaborations.” *Human Computer Interaction*, **3**(1), 31-58, Reprinted in Greif, 1988.

Lai, K.-Y. and Malone, T. W. (1988) “Object Lens: A 'spreadsheet' for cooperative work.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 115-124, Portland, Oregon, September 26-28, ACM Press.

Lantz, K. (1986) “An experiment in integrating multimedia conferencing.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, Austin, Texas, December, ACM Press. Reprinted in Greif, 1988.

Lederberg, J. (1989) *Excitement and fascination of Science.* **3** Palo Alto, California, Annual Reviews, Inc,

Lederberg, J. and Uncapher, K. (1989) “Towards a national collaboratory.” The Rockefeller University, March 17-18. Also available from the National Science Foundation.

Leland, M. D. P., Fish, R. S. and Kraut, R. E. (1988) “Collaborative document production using Quilt.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 206-215, Portland, Oregon, September 26-28, ACM Press.

Mackay, W. (1988) “More than just a communication system: Diversity in the use of electronic mail.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 344-353, Portland, Oregon, September 26-28, ACM Press.

Malone, T. W., Grant, K. R., Lai, K.-Y., Rao, R. and Rosenblitt, D. (1987) “Semi-structured messages are surprisingly useful for computer-supported coordination.” *ACM Transactions on Office Information Systems*, **5**(2), 115-131, April. Reprinted in Greif, 1988.

Malone, T. W., Grant, K. R. and Turbak, F. A. (1986) “The Information Lens: An Intelligent system for information sharing in organizations.” In *Proceeding of the SIGCHI Human Factors in Computing Systems*, 1-8, Boston, Mass, April 13-17, Association for Computing Machinery.

Mayer, M. (1977) “The telephone and the uses of time.” In *The Social Impact of the Telephone*, edition). Cambridge, Mass, Pool, I., MIT Press.

MIT (1983a) “Talking heads.” In *Discursions*, edition). Boston, Mass, Architecture Machine Group, MIT. Optical disc.

MIT (1983b) “Zero bandwidth video.” In *Discursions*, edition). Boston, Mass, Architecture Machine Group, MIT. Optical disc.

Nelson, T. (1987) *Computer lib/Dream machines.* Redmond, Washington, Tempus Books,

Pliskin, N. (1989) “Interacting with electronic mail can be a dream or a nightmare: a user's point of view.” *Interacting with Computers*, **1**(3), 259-272, December.

Quigley, A. (1989) “The politics of access: Identifying components of a research agenda.” *Research in education*, **1**(1), 3-5, April.

Root, W. R. (1988) “Design of a multi-media vehicle for social browsing.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 25-38, Portland, Oregon, September 26-28, ACM Press.

Sarin, S. and Greif, I. (1985) “Computer-based real-time conferencing systems.” *IEEE Computer*, **18**(10), 33-45, Reprinted in Greif, 1988.

Scott, P. B. (1984) *The robotics revolution.* Oxford, UK, Basil Blackwell Publisher,

Smith, R. B. (1988) “A prototype futuristic technology for distance education (working draft).” In *NATO Research Workshop on New Directions in Education Technology*, Cranfield, England, November,

Sproull, L. and Kiesler, S. (1986) “Reducing social context cues: Electronic mail in organizational communication.” *Management Science*, **32**(11), 1492-1512, Reprinted in Greif, 1988.

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S. and Tatar, D. (1987a) “WYSIWIS revised: Early experiences with multiuser interfaces.” *ACM Transactions on Office Information Systems*, **5**(2), 147-167, April.

Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. and Suchman, L. (1987b) “Beyond the chalkboard: Computer support for collaboration and problem solving in meetings.” *Communications of the ACM*, **30**(1), 32-47, Reprinted in Greif, 1988.

Stevens, C. H. (1986) “Electronic organization and expert networks: Beyond electronic mail and computer conferencing.” Working paper 90s:86-021, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts, May.

Stults, R. (1988) “Experimental uses of video to support design activities.” Xerox PARC, Palo Alto, California, December.

Thomas, R. H., Forsdick, H. C., Crowley, T. R., Schaaf, R. W., Tomlinson, R. S., Travers, V. M. and Robertson, G. G. (1985) “Diamond: A Multimedia Message System Built on a Distributed Architecture.” *IEEE Computer*, **18**(12), 65-78, Reprinted in Greif, 1988.

Trigg, R. H. (1988) “Guided tours and tabletops: Tools for communicating in a hypertext environment.” In *Proceedings of the Conference on Computer-Supported Cooperative Work*, 216-226, Portland, Oregon, September 26-28, ACM Press. Reprinted in Greif, 1988.

Watson, R. T., DeSanctis, G. and Poole, M. S. (1988) “Using a GDSS to facilitate group consensuss: Some intended and unintended consequences.” *MIS Quarterly*, **12**(3), 463-477, September.

Winograd, T. (1988a) “A language/action perspective on the design of cooperative work.” *Human Computer Interaction*, **3**(1), 3-30, Reprinted in Greif, 1988.

Winograd, T. (1988b) “Where the action is.” *Byte Magazine*, , December.

Winograd, T. and Flores, F. (1986) *Understanding computers and cognition: A new foundation for design.* Norwood, New Jersey, Ablex,