Computer Support for Real Time Collaborative Work

Saul Greenberg and Ernest Chang

Advanced Technologies
Alberta Research Council
6815 8 St NE
Calgary, Alberta, Canada T2E 7H7
phone: (403) 297-2674
email: saul@noah.arc.cdn

Abstract—A large portion of an office worker's time involves real time collaboration with fellow associates. Many traditional tools assist this process, such as telephones, meeting rooms, specialized media such as whiteboards, and so on. This paper is concerned with the potential of the computer as a tool to further enhance the group work process through direct support of real time communication needs and the specific collaboration requirements of the group. General computer support for four styles of real-time interactions are distinguished and surveyed: 1) face to face meetings; 2) remote conferencing; 3) casual real time interaction; and 4) multi-user applications. Each topic is introduced, motivations discussed, and the key technical systems and related research described.

1 Introduction

Although computers are now familiar tools used by people to pursue their own individual tasks, they have not, until recently, been exploited to assist people working together. Computer-supported cooperative work (CSCW) is a new multi-disciplinary research field exploring the potential of the computer to support group work. CSCW acts as an umbrella to a variety of specializations including electronic mail, asynchronous conferencing, bulletin boards, group decision support systems, collaborative authoring, group knowledge elicitation, shared workspaces, and so on.

This document will focus only on those CSCW systems that support real time communication and collaboration between people. General computer support for four styles of real-time interactions are distinguished and surveyed: 1) face to face meetings; 2) remote conferencing; 3) casual real time interaction; and 4) multi-user applications.

The paper begins with a discussion of *face-to-face formal/semi-formal* meetings, where all participants are located in the same meeting room and have ready access to the same physical materials. Formal meetings are typically scheduled, are goal-oriented, have invited participants, and often involve an organizer or chairperson. The discussion of computer support revolves around the computer as both a facilitator to the meeting and as an extended media tool.
The subsequent section considers *remote conferencing*, where some or all attendees at a real-time meeting are located at different sites. The role of technology in remote conferencing emphasises bringing people and their materials together in a way that allows participants to orchestrate the conversation as effectively as they do in face-to-face meetings, and to share their otherwise inaccessible materials over distance.

The discussion continues on the theme of computer-support for *casual real-time interactions*. Spontaneous unplanned meetings are often crucial for bringing together people for collaborations, yet the opportunity for accidental encounters has an exponential decay with distance. Computer mediation is explored as a way of bringing distance-separated people into contact through frequent, unplanned, high-quality, and real-time interactions.

While the above distinctions derive from social interaction, another approach stems from *multi-user applications.* These are those problem-specific programs that support the distributed and often persistent interactions of a group of persons who work together through the roles intrinsic to the task at hand. These applications tend to be structured less as meetings and more as integrations of distinct roles in large endeavours, such as the teamwork of captain, navigator, pilot and engineer on a ship.

The paper closes with a list of the relevant books, journals, and technical conferences related to computer supported cooperative work.

Cite as:
Greenberg, S. and Chang, E. (1990) "Computer support for real time collaborative work." *Congressus Numerantium*, **75**, pp. 247-262, June.
First published in Proceedings of the Conference on Numerical Mathematics and Computing, Sept 28-30, Winnipeg, Manitoba, 1989

2 Computer-support in face to face meeting rooms

Face to face meetings occur when all participants are located in the same room, perhaps sharing a common work area. This section elaborates on the role of conventional and computerized media in the meeting room, and the effects the computer may have on the group facilitative process.

2.1 Conventional media. Participants at formal or semi-formal face to face meetings often use or are exposed to a variety of conventional media tools. These usually involve some large communal work surface and private work areas that may be made public later. Example media tools include overheads, slides, videos, whiteboards, flipcharts, pointers, coloured pens, private notepads. The diversity of tools support different meeting and participant needs, needs that often include

- presentation of prepared material,
- brainstorming,
- interactive design,
- private and group note-taking,
- private and group annotation,
- gesturing.

Stefik *et al* (1987b) observes that “media influence the course of a meeting because they interact strongly with participants’ resources for communication and memory.” As a group resource, the choice of media and its capabilities becomes crucial. Consider Tang’s (1989) comments when describing how different materials can affect an interactive design meeting.

> The [design] workspace can be composed of: notebooks, whiteboards, shared sheets of paper on a table, tracing paper, networked computer workstations, and combinations thereof. These diverse options represent a variety of capabilities: portable, eraseable, simultaneously accessible, layerable, editable. The composition of the workspace determines what capabilities are afforded, and thus plays a major role in structuring the participants’ workspace activity (Tang, 1989 page 66).

Existing media aids are rather limited. Consider, for example, the pervasive whiteboard and its equivalents. Although it is an excellent and quite general medium for providing a shared and focused memory for a meeting, a whiteboard is physically quite restricted (Stefik *et al*, 1987b). The amount of surface available for drawing is fixed. Only a few people can simultaneously use it. Whiteboards are ill-suited for re-arranging existing items. They cannot normally keep a record of the artifacts drawn, nor do they allow previously prepared information to be imported easily.

Even a single media form can have different capabilities. Tang, for example, describes two design teams using large sheets of paper as a shared workspace. One group tacked the paper on top of a whiteboard, the other put it on a table. He noticed that the different location of the media had a profound effect on how the group used them (Tang, 1989).

Orientation. When people sat around the table, drawings made on the table-mounted paper were oriented in different directions. Although people had greater difficulty drawing and perceiving the images, orientation proved a resource for facilitating the meeting. Because drawings faced a particular person, a context and an audience was established. Marks made by participants that were aligned to an image conveyed support and focus. People working on their own image used orientation as a “privacy” boundary until they were ready to call in the group’s attention. The whiteboard mounted paper had none of these features.

Proximity. Tang noticed that when participants were huddled around the tabletop workspace, the workspace played a key role in mediating the conversation. This role was lessened in the whiteboard situation where people were seated several feet from it.

Simultaneous Access. Given good proximity, a high percentage (45—68%) of workspace activity around the tabletop involved simultaneous access to the space by more than one person. Although parallel activity decreases bottlenecks, it also decreases focused interaction.

Another problem with conventional media is the lack of integration between the available forms, and the technically difficult problems of using more than one media in a meeting. Furthermore, since media aids are usually designed to be general purpose, their use in task-specific meeting activities is often a compromise.

2.2 Computerized media. Given the trade off between capabilities and limitations of physical media, the computer has potential to include the best of existing media tools while limiting its physical restrictions, and the potential to extend what media currently enables for meetings. Consider the incomplete list of benefits below.

As a presentation tool:

- modifiable on-line still graphics can replace both the overhead and slide projector;
- running applications can be projected;
- video recorders and optical disks can be controlled through software.

As a tool to encourage participant interaction:
- participants can rapidly access and simultaneously manipulate the objects in the group workspace without leaving their seats;
- they can publicly refer to and annotate the workspace by (say) a public pointer and pen;
- they can import their own on line material to the workspace;
- they can save and further manipulate the public workspace in their own private area.

As a specialized meeting tool:
- using existing equipment, a meeting can bring in a particular flavour of meeting tool designed to facilitate a special meeting task, such as brainstorming, decision support, and so on.

As a facilitator:
- participants can vote anonymously, with results tabulated automatically;
- a textual sub channel between participants could allow private conversations between participants;
- social and emotional feedback can be entered and tabulated by the machine to indicate the current mood of the group.

2.3 Technical advances. Xerox PARC pursued the idea of computer support for face-to-face meetings in an experimental meeting room for small groups known as *CoLab* (Stefik *et al*, 1987b). The room is arranged with one workstation per participant, as well as a very large touch-sensitive screen and stand-up keyboard. 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 and Cognoter are discussed below in their respective roles as an extension to an existing media device, and as a new tool for augmenting meeting activity.

Like a chalkboard, the Boardnoter supports informal freestyle sketching and erasing on a communal work area by participants. Unlike a chalkboard, the Boardnoter allows:

- 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 *telepointer*[1] (Stefik *et al*, 1987a).

The Cognoter is quite a different and more formal 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.

Another well-known computer-supported face-to-face meeting room is *Project Nick.* Like CoLab, the room contains a workstation per participant and a large shared screen (Cook *et al*, 1987; Ellis *et al*, 1988a). Nick's mandate is to apply automated facilities to the process, conduct, and semantic capture of design meetings (Cook *et al,* 1987). Its prototype facilities include:

- group worksurfaces visible to all participants;
- a "facilitation" system that keeps members attuned to the meeting by displaying on the group worksurface the meeting's progress and results;
- software to display agenda information;
- software to permit entry of information into a persistent "group memory";
- subchannels for textual communication between two or more meeting participants, including:
  - public information sent from one participant to all others;
  - private information sent between a subset of participants;
  - binary information for voting;
  - pre-formatted mood indicator messages sent by a single keypress to the meeting chairperson;
- a group editor for creating and editing lists and outlines (Ellis *et al,* 1988b);
- software for capturing and processing meeting statistics to help quantify the meeting effectiveness.

2.4 Socio-emotive concerns. Computer support for face-to-face meetings involves much more than a clever bag of software tricks. It is too easy to concentrate on the rational communication channels of a meeting. Just

[1]A *telepointer* is an individual's cursor made visible to all participants viewing the shared image.

as important are the affective channels that include social, emotional, and organizational information (Rein and Ellis, 1989).

Consider the *Capture Lab*, a computer-supported meeting room whose construction emphasised the need for careful design of all aspects of the room (Mantei, 1988; EDS, 1988). The subtle effects of seemingly trivial items such as seating, viewing distances between participants, availability of a front screen, and access protocols had a profound effect on the way the computer-supported meeting was run. Despite trying to make the new meeting environment as close as possible to the old environment the participants were used to, Mantei noticed quite dramatic changes in the group's dynamics. These include:

- changes in seating arrangement;
- shifts of the "power position" at the conference table;
- a loss of privacy for individual participants;
- changes in protocols for designating the scribe (the person recording the meetings progress on the electronic blackboard);
- dominance of the screen as the group's focus of attention.

Because the electronic blackboard became a large part of the meeting, senior management, for example, vied for the "power seat" where their field of vision included not only the blackboard, but also all the participants and the main door to the room as well. Less senior participants would adjust themselves accordingly.

Another example of the importance of the affective channel was reported in Rein and Ellis' (1989) study of several software design teams using the Project Nick meeting room. The subchannel messaging facility that allowed people to communicate to each other textually emerged as an effective way for quiet people to influence the meeting. Rein and Ellis relate several episodes where the verbal communication channel was getting off track due to heated argument and long diatribes. The subchannel was used by participants to either tell the chairperson to regain control of the meeting, or to communicate concerns to other participants. In one episode, an extensive subchannel discussion occurred, effectively by-passing the current verbal speaker who was monopolizing the floor.

2.5 Limitations and unknowns. Although several "success" stories of computer-supported meeting rooms exist, it is still too early to judge their general effectiveness. 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.

However, computer support should do *at least* as well as conventional media, simply by replacing that media without altering functionality. The real potential exists because computer support could surpass conventional media because it does not share their inherent physically restrictions.

On the other hand, computer support, done poorly, can bring a meeting to its knees. For example, availability of a computer console in a meeting makes it easier for a participant to import and get distracted by his private work, possibly diffusing the effectiveness of the gathering. Just having terminals/keyboards/mice at the conference table is a distraction that may splinter, rather than unify, the meeting's members.

Perhaps designers of current meeting tools are too ambitious in their scope, for they often strive to create very general meeting tools. It may be more beneficial to design quite specialized tools for supporting specific meeting facets. In this vein, Ernie Chang of the Alberta Research Council is taking a general shared workspace tool and moulding it into a presentation tool that allows a presenter to hand off gesturing and annotation control to selected members of the audience. Given this and other specialized foci, we can observe those meeting aspects and base the tool design on hard observations rather than intuition. As Tang (1989) writes:

> …[we] need to understand what participants actually do in a [meeting] activity in order to guide the development of technology (especially advanced computer tools) to support this activity.
>
> (Tang, 1989 page iii).

3 Computer support for remote conferencing

Remote conferencing brings participants together in formal or semi-formal meetings, 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; same-time, different place meetings; and so on.

Two aspects of remote conferencing are distinguished here: *tele-presence* and *tele-data*.

3.1 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; 1989). Each participant has the option of displaying one or more windows representing the other participants. When text is typed by one or more "speakers," it is broadcast to all participants and displayed in the window representing that speaker. 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 telepresence 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. 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). In both systems, the feeling of presence is striking. Participants can effectively orchestrate the conversation through natural eye contact and head movement.

3.2 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 present or access physical 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, 1989). 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 telepointer). 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 (Lantz, 1986; Gust, 1989; Ensor *et al*, 1988; Ensor, 1989). At the Alberta Research Council, Saul 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. 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 (1989).

3.3 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 which 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].

3.4 Limitations and unknowns. Although both tele-data and tele-presence are clearly important, their effective implementation and the role they play in particular types of meetings are not well known. 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

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

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.

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). 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, microphones, wires, cameras and monitors may be a significant intrusion to the meeting, particularly to those participants who do not feel comfortable with the technology.

Given that 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 multidisciplinary teams of sociologists, anthropologists and computer scientists to study how people communicate through a video channel. In one of their experiments, they observe a team of three architects working together solely through a video system that transmits images of both the person and his data (Stults, 1989; Harrison, 1989). The knowledge acquired through this process will help determine the design requirements behind truly useful remote conferencing systems.

4 Casual real-time interactions

Computer support for both remote and face-to-face conferencing has so far addressed 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. Kraut, Egido and Galegher (1988) argue that many interactions are required for people to find partners for collaborative work. They especially emphasise the importance of brief unplanned encounters where bits of technical and personal information are exchanged "on the fly" (Root, 1988).

Yet the bottleneck to rich spontaneous interactions is distance. As Kraut et al (1988) report, the communication contacts between people is well-proven to have an exponential decay with distance, leading to a decrease in potential collaborations. The number of collaborations, for example, drops off sharply when one contrasts people working on the same floor, on different floors, and in different buildings.

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 below: *video hallways* and *shared alternate reality*.

4.1 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, 1989; 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 (1987) reported that 70% of all Video Wall interactions were spontaneous, and the other 30% planned. 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. Users also reported dissatisfaction with the poor image quality of slow scan video.

A second prototype video hallway is *Cruisin'* (Root, 1988; Fish, 1989). While Xerox's Video Wall directly connects two physical locations, Cruisin' attempts to create a virtual community where everyone has instance access to everyone else. Cruisin' 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 Cruisin' 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 Cruisin'? 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 Cruisin' 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 (1988) 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. 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. Cruisin' 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 Cruisin' is very much a prototype, and much remains to be built and tested. In contrast, the TeleCollaboration project of Corey *et al* (1989) supports a high-speed full video link between two *US West* sites in Colorado separated by a distance of 100 kilometres. Their system allows users to "walk" through the physical hallways and offices of distant sites. They can even scan larger rooms through 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.

4.2 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 intensional 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 "teleporter." 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 traveling 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.

4.3 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 unlikely that most institutions will seriously consider video hallways. 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).

4.4 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 n seconds (currently set at 60). 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.

5 Multi-user applications

Computer-supported collaborative work through the conferencing or meeting paradigm focuses on social interaction as the primary vehicle through which work is conducted. Multi-user applications, on the other hand, focus on a specific task and its representation as the central problem and considers individuals as agents through which various components of the task are achieved.

Software supporting the conference paradigm typically address and implement social mechanisms such as floor control and focus of attention. Access to computer-based applications or data are generally on a "let's all see it and let one person at a time operate it" basis.

Multi-user applications are supported by domain-specific software that address a particular problem and integrate the several persons that may be involved through task-specific roles and behaviours, perhaps with a participant being made aware of anothers activities only when required. In coordinating the actions of the captain, navigator, pilot and engineer of a ship, the software may have to handle input from each at different levels of priority with different routines, provide custom views as well as shared views of a common data base, provide locking and conflict-resolution for writable objects, provide multi-media (voice, text and graphics) communication between the agents and give access to various data and functions through different levels of security and privilege.

The mechanisms that support computer conferencing and the meeting paradigm are based on generic social concerns which may not be relevant to specific multi-user applications. On the other hand, a particular problem involving several persons may require coordination mechanisms which are not seen in any other domain. Research in the former aims to provide insights and techniques of wide applicability; work in task-specific multi-user applications give immediate gratification and delays generalization until more successful cases have been understood. Both approaches are important in the new field of computer supported collaborative work, and reflect the richness and diversity of this emerging area.

6 Further readings and sources

This section is a guide for people interested in current and future literature in computer-supported cooperative work. Due to the youth and diversity of CSCW, there are few specialized books or journals available. Most of the literature is scattered amongst many journals, proceedings of conferences and workshops, and technical reports.

The meagre list below indicates a) books and journals that cover the CSCW area, and b) mainstream and related conferences and workshops on CSCW.

6.1 Books and journals. Perhaps the best overview of the CSCW discipline is provided by Irene Greif (1988). Her collected readings span from the earliest visions of CSCW to present day theory and practice. This book is required reading for anyone interested in the field.

There are no dedicated journals for CSCW. Still, the ones listed below do publish occasional papers of interest.

- "ACM Transactions on Office Information Systems" is probably the best bet. Volumes 5(2) 1987, 6(3) 1988 and 6(4) 1988 were special CSCW issues.
- The mainstream human-computer interface journals have occasional articles on CSCW. Journals include
  - International Journal of Man Machine Studies;
  - Human Computer Interaction;
  - Behaviour and Information Technology; and
  - Interacting with Computers.
- ACM's "SIGCHI Bulletin" and "SIGOIS Bulletin" occasionally contain relevant articles and summaries of conferences and workshops.
- The odd high-quality article has appeared in "Communications of the ACM" and "IEEE Computer."
- The December 1988 issue of "Byte" had a special section devoted to CSCW papers.

6.2 Conferences and workshops. The major conference for CSCW is the bi-annual "ACM Conference on Computer-Supported Cooperative Work." The first conference was held in 1986 (Austin, Texas), the second in 1988 (Portland, Oregon), and the third is scheduled for October 1990 in Los Angeles. The proceedings are of very high quality, with the collected papers covering most contemporary work. While the 1986 proceedings is difficult to obtain, the 1988 one is available from ACM Press.

There have also been several workshops on CSCW. Unfortunately, their proceedings are also hard to come by. If you can get it, the collected position papers of the "1989 Groupware Workshop" (Palo Alto) provide a flavour of what implementors of CSCW systems are working on, and the issues they are being confronted with.

Several other conferences and their proceedings are worthy of mention.

- The "First European Community Conference on Computer Supported Cooperative Work" (EC-CSCW) was held on September 13—15 in Gatwick, London, UK.
- ACM's "SIGCHI Conference on Computer-Human Interaction" and "SIGOIS Conference on Office Information Systems" sometimes have special sessions, panels or papers on CSCW.
- The University of Guelph sponsors the "Symposium on Computer Conferencing". Most papers in the second symposium were centred around distance education issues (University of Guelph, 1987).
- The "2nd IEEE ComSoc International Workshop on Multimedia Communications" was held in 1989. Although this particular workshop emphasised the network demands of multi-media communications, there were several CSCW sessions.

7 Bibliography

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.* Ed. M. Taylor et. al. North Holland 1989.

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

Clark, W. (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.

Cook, P., Ellis, C., Graf, M., Rein, G., and Smith, T. (1987) "Project Nick: Meetings augmentation and analysis." *ACM Transactions on Office Information Systems*, 5(2):132--146, April.

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

EDS (1988) "The Capture Lab." Videotape produced by the EDS Centre for Machine Intelligence, Ann Arbour, Michigan.

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*, pages 13--24, Portland, Oregon, September 26--28.

Ellis, C., Gibbs, S., and Rein, G. (1988a) "The Groupware project: An overview." Technical report STP-033-88, MCC, Austin, Texas, January.

Ellis, C., Gibbs, S., and Rein, G. (1988b) "Design and use of a group editor." Technical report STP-263-88, MCC, Austin, Texas, September.

Engelbart, D. and English, W. (1968) "A research center for augmenting human intellect." In *Proceedings of the Fall Joint Computing Conference*, volume 33, pages 395--410. AFIPS Press, Montvale, NY, Fall.

Ensor, J. (1989) "Rapport: A multimedia conferencing system." In the *ACM SIGGRAPH Video Review Supplement to Computer Graphics*, Issue 45(5), ACM Press, Baltimore, MD. Videotape.

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

Farallon (1988) "Timbuktu user's guide." Farallon Computing Inc., Berkely, California.

Fish, R. (1989) "Cruiser: A multi-media system for social browsing." In the *ACM SIGGRAPH Video Review Supplement to Computer Graphics*, Issue 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*, pages 7--15, Austin, Texas, December 3--5.

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

Greenberg, S. (1989) "Sharing views and interactions with single-user applications." Research report, Advanced Technologies Department, Alberta Research Council, Calgary, Alberta, Canada, June.

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

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." In the *ACM SIGGRAPH Video Review Supplement to Computer Graphics*, Issue 45(4), ACM Press, Baltimore, MD. Videotape.

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

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

Lantz, K. (1986) "An experiment in integrating multimedia conferencing." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, Austin, Texas, December.

Mantei, M. (1988) "Capturing the Capture concepts: A case study in the design of computer-supported meeting environments." In *Proceedings of the Conference on Computer-Supported Cooperative*

Work, pages 257--270, Portland, Oregon, September 26--28.

MIT (1983a) "Talking heads." In *Discursions*, Architecture Machine Group, MIT, Boston. Optical disc.

MIT (1983b) "Zero bandwidth video." In *Discursions*, Architecture Machine Group, MIT, Boston. Optical disc.

Rein, G. and Ellis, C. (1989) "The Nick experiment reinterpreted: Implications for developers and evaluators of groupware." *Office: Technology and People*, 5(1).

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

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

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

Stefik, M., Bobrow, D., Foster, G., Lanning, S., and Tatar, D. (1987a) "WYSIWIS revised: Early experiences with multiuser interfaces." A*CM 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.

Stults, R. (1989) "Experimental uses of video to support design activities." Research report SSL-89-19, Xerox PARC, California.

Tang, J. (1989) "Listing, drawing, and gesturing in design: A study of the use of shared workspaces by design teams." Phd thesis, Department of Mechanical Engineering, Stanford University, California, April. Also available as research report SSL-89-3, Xerox PARC, California.

University of Guelph (1987) "The Second Guelph Symposium on Computer Conferencing." University of Guelph, Guelph, Ontario, Canada, June 1--4.