

Human and Technical Factors of Distributed Group Drawing Tools

**Saul Greenberg
Mark Roseman
David Webster
Ralph Bohnet**

1992

Cite as:

Greenberg, S., Roseman, M., Webster, D. and Bohnet, R. (1992) "Human and technical factors of distributed group drawing tools." *Interacting with Computers*, **4**(1), pp. 364-392, December. Butterworth-Heinemann. (Special edition on CSCW, Tom Rodden ed.)

This manuscript was used to produce the paper cited above. However, they are not identical. Small word-level differences exist. Additionally, Figure 1 is a poor quality copy of the actual figure, and Figure 2 is not available.

Human and Technical Factors of Distributed Group Drawing Tools

Saul Greenberg
Mark Roseman
Dave Webster

Department of Computer Science, University of Calgary, Calgary, Canada T2N 1N4

Ralph Bohnet
MPR TelTech Ltd, 8999 Nelson Way, Vancouver, Canada

Abstract

Groupware designers are now developing multi-user equivalents of popular paint and draw applications. Their job is not an easy one. First, human factors issues peculiar to group interaction appear that, if ignored, seriously limit the usability of the group tool. Second, implementation is fraught with considerable technical hurdles. This paper describes the human and technical factors that have been met and handled by researchers and implementors of group drawing tools. We emphasize our own experiences building four systems supporting remote real time group interaction: GROUPSKETCH and XGROUPSKETCH, both multi-user sketchpads; GROUPDRAW, a prototype object-based multi-user drawing package, and GROUPKIT, a groupware toolkit. On the human factors side, we summarize empirically-derived design principles that we believe are critical to building useful and usable collaborative drawing tools. On the implementation side, we describe our experiences with replicated versus centralized architectures, schemes for participant registration, multiple cursors, network requirements, and the structure of the drawing primitives. A brief survey of other approaches to group drawing is included as well.

Keywords: Group drawing, shared workspace, real time remote conferencing, desktop conferencing, computer supported cooperative work, groupware.

1: Introduction

Most research efforts in geographically distributed conferencing have been in the field of *tele-presence*—a way of giving distributed participants a feeling that they are in the same meeting room (Egido 1988; Johansen and Bullen 1984; MIT 1983). The goal of tele-presence is to transmit both the explicit and subtle dynamics that occur between participants. These include body language, hand gestures, eye contact, meta-level communication cues, knowing who is speaking and who is listening, voice cues,

focusing attention, 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. *Tele-data*, on the other hand, allows participants at a meeting to present or access physical materials that would normally be inaccessible to the distributed group (Greenberg and Chang 1989). 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 centered around a computer workstation, the networked computer can become a valuable medium for people to share on-line work with each other.

In this document, we will focus on tele-data that provides small groups (2 to ~8 people) with real-time access to a shared drawing space via multi-user equivalents of the now-common paint and draw programs. We describe both human and technical factors behind our design and implementation of four systems supporting real time group interaction.

GROUPSKETCH is a minimalist multi-user sketchpad that takes over the entire computer display (Figure 1) (Greenberg and Bohnet 1991). Its users all see exactly the same thing, and anyone can draw, type, erase, move their cursors around the drawing, and save or restore the image. It is a robust program that we have made available to the research community via anonymous ftp.

XGROUPSKETCH is a re-implementation of GROUPSKETCH (Figure 2). It differs in that it runs within a standard re-sizeable X window, has a scrollable drawing surface, and includes features such as different pen sizes and colors, font types and font sizes, and the ability to save and restore multiple screen images. It too is a moderately robust program.

GROUPDRAW is an object-oriented drawing program with features similar to those of most structured drawing packages (Figure 3) (Greenberg, Roseman, Webster and Bohnet 1992). Unlike the sketchpads, users create objects (such as lines and squares) that can then be manipulated. GROUPDRAW is a prototype developed as a throw-away experimental platform; it is buggy and fragile, and will not be released to the community.

GROUPKIT is a toolkit for building distributed real time groupware (Roseman and Greenberg 1992a & 1992b). Even though GROUPKIT applications can go far beyond drawing, the toolkit design is based upon many of the lessons we have learnt from the applications described above. GROUPKIT is now under development, and will be released to the community on completion.

These implementations are not unique. The technological breakthrough of building computational group drawing spaces was reported many years ago (Sarin and Greif 1985), and other researchers are busy constructing their own group drawing packages and toolkits (see final section). However, we believe the time is ripe to proceed beyond mere replication and move towards empiricism, where lessons are drawn

from experience and formulated as empirical design rules (Gaines, 1991). Our intent in this paper is to use our systems, concentrating on GROUPSKETCH and GROUPDRAW, to highlight human factors issues critical to the design of real time collaborative drawing tools, and to pass on the important technical factors we have encountered building such systems.

We begin the paper with a literature summary of human factors studies of face to face design teams and the resulting design principles generated from them. We then explain how these principles were incorporated into our groupware designs, and the early experiences people had using these systems. The following section describes the technical factors encountered during implementation, concentrating on the choice of a replicated over a centralized architecture, handling of participant registration, displaying of multiple cursors, the network requirements, and the underlying structure of the drawing primitives. We conclude with a brief survey and perspective of other group drawing systems.

2: Designing for the human factors of small group meetings

Almost every group meeting includes a process where participants express, discuss, and develop ideas. It is a creative forum where people are encouraged to present their thoughts to the group for feedback, to build upon the ideas presented by fellow members, and to problem-solve. Participants typically use some large communal *work surface*—a group drawing area—to facilitate their interactions. Typical media now used include whiteboards, flipcharts, large sheets of paper, as well as a variety of colored pens for drawing.

Our aim is to apply existing human factors knowledge of face to face meetings to the design of workstation conferencing tools supporting remote work surfaces. This section will describe how people use conventional work surfaces, the implications for the design of a workstation-based work surface, and how our systems instantiated the design recommendations. While most of the human factors knowledge is derived from studies of small group design teams brainstorming through initial design ideas, we are convinced that the findings are generally applicable to most uses of a shared work surface.

2.1: Understand collaboration.

In order to design a software-based distributed work surface, we must have an adequate understanding of how traditional ones are used in face to face group meetings. Indeed, Grudin has identified a lack of understanding of group behavior to be one of the reasons why groupware has not been generally successful (Grudin 1989). He asserts that designers rely too heavily upon their own intuition, which is often based upon experiences that may not be applicable to the target group as a whole.

For example, an intuitive “conventional” view of the communal work surface would consider it merely as a medium for creating and storing a drawing artifact (Tang 1989). Bly (1988) disproves this naive view. She studied two designers communicating through three different media offering different access to a drawing surface: face to face including a shared sketchpad; over a video link that included a view of the other person and their personal drawing surface; and over the telephone. From her observations, she asserts that the drawing process—the actions, uses, and interactions on the drawing surface—are as important to the effectiveness of the collaboration as the final artifact produced. Bly also noticed that allowing designers to share drawing space activities increases their attention and involvement in the design task. When interaction over the drawing surface is reduced, the quality of the collaboration decreases.

Tang refined Bly’s findings even further through his ethnographic study of eight short small-team design sessions (Tang 1989 & 1991; Tang and Leifer 1988; also see Garner, Scrivner, Clarke et al, 1991 for further work in this area). Each team used large sheets of paper as a shared work surface and were given problems to solve. Some teams placed the paper on a table, others tacked it to a whiteboard. Even this simple difference had a profound effect on how the group used the shared work surface. Several important observations resulted:

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 group using whiteboard mounted paper did not exhibit these behaviors.

Proximity. Tang noticed that when participants were huddled around the table-mounted paper, the sketchpad played a key role in mediating the conversation. This role was lessened in the whiteboard situation where people were seated several feet away.

Simultaneous Access. Given good proximity, a high percentage (45–68%) of work surface activity around the tabletop involved simultaneous access to the space by more than one person.

Tang (1989) built a descriptive framework to help organize the study of work surface activity, where every user activity was categorized according to what action and function it accomplished, as listed below.

Actions:

- *listing* produces alpha-numeric notes that are spatially independent of the drawing;

- *drawing* produces graphical objects, typically a 2-dimensional sketch with textual annotations that are attached to the graphic;
- *gesturing* is a purposeful body movement that communicates specific information, eg pointing to an existing drawing.

Functions:

- *storing information* refers to preserving group information in some form for later recall;
- *expressing ideas* involves interactively creating representations of ideas in some tangible form, usually to encourage a group response;
- *mediating interaction* facilitates the collaboration of the group, and includes turn-taking and focusing attention.

Tang’s classification of small group activities within this framework revealed that the “conventional” view of work surface activity—storing information by listing and drawing—constitutes only ~25% of all work surface activities. Expressing ideas and mediating interaction comprised the additional ~50% and ~25% respectively. Gesturing, which is often overlooked as a work surface activity, played a prominent role in all work surface actions (~35% of all actions). For example, participants enacted ideas by using gestures to express them, and gestures were used to signal turn-taking and to focus the attention of the group.

2.2: Implications for design of a work surface

Tang’s observations led him to derive six design criteria that shared work surface tools should support. He stresses the importance of allowing people to gesture to each other over the work surface, and emphasizes that the process of creating a drawing is in itself a gesture that must be shown to all participants through continuous, fine-grained feedback. Another key point is that the tool must not only support simultaneous activity, but also encourage it by giving participants a common view of the work surface. The six design criteria plus a summary of their rationale are listed in detail in the first two columns of Table 1 (condensed from Tang 1989), and deserve special attention of the reader. These criteria form the foundation for the rest of this paper.

—Table 1 around here—

2.3: The GROUPSKETCH interface

GROUPSKETCH is a simple group sketching tool that allows an arbitrary number of people to draw on a virtual piece of paper (the screen) using the mouse (Greenberg and Bohnet 1991). It is designed around the criteria listed in Table 1. Its main features are:

- it uses a strict what you see is what I see (WYSIWIS) display;
- multiple, active cursors that identify their owners are always visible on all displays;
- simultaneous interaction is fully supported, and any user can perform any action at any time;
- any user action (cursor movement or drawing), no matter how small, is immediately visible on all screens
- drawing, gesturing and listing are as modeless as possible.

Column 3 of Table 1 provides detail of the GROUPSKETCH user interface features and how they were designed around Tang's six criteria.

Figure 1 displays a typical GROUPSKETCH screen with four participants engaged in a design session. On the left is the shared work surface where people draw, enter text, or gesture. Every person has a cursor labelled with their name. All participants see the same work surface on their display, and every movement of the cursor and change in the drawing is immediately visible on all displays. Each participant is represented by a unique labelled caricature located outside the work surface on the right of the screen. While audio is not directly supported, we expect a full duplex audio channel to be available by other means (eg conference calls).

—Figure 1 around here—

Four action modes are supported: gesturing through cursors, drawing, typing, and erasing (Figure 1). With no mouse buttons or keyboard keys pressed, the cursor portrays the image of a pointing hand (Sandy's cursor). To draw free-style, the user depresses the left mouse button of a three-button mouse, changing the cursor from a hand to a pen (Saul's cursor). The pen-shaped cursor also appears automatically when typing. Pressing the middle mouse button changes the cursor into a large arrow to draw participants' attention (Irene's cursor). Users can erase graphics or text in the work surface by holding down the right mouse button, which changes the shape of the cursor into an eraser (Wilf's cursor).

The menu on the right of Figure 1 allows a person to privately save an image, retrieve a previously stored image to the group display, clear the public work surface, or leave the collaboration (leaving other participants in the meeting). Menu selections and cursor movements outside the work surface are private and are not broadcast to other workstations (this is the only part of GROUPSKETCH that is not WYSIWIS). Loading an image or clearing the work surface affects all participants' screens.

XGROUPSKETCH is similar to GROUPSKETCH with several important differences. First, it runs within a standard X window (Figure 2), and users can manipulate this window in all the usual ways. Second, it is functionally richer. Through a control panel with pull down menus, users can choose different line styles, line colors, and fonts. Multiple images can be saved as pages in files through the file menu. A user can raise a list of all XGROUPSKETCH participants, shown in a separate window, through the “Users” button. Feedback is located on the top right, with a panel specifying the font, line width, the number of users, and arbitrary system messages. Third, the window is a viewport into a large drawing area, and users can scroll independently to and work on any part of the drawing they wish.

—Figure 2 around here—

We have performed limited usability studies on GROUPSKETCH under relatively informal conditions (Greenberg and Bohnet 1991)¹. In a typical GROUPSKETCH scenario, participants converse and interact as they would over a shared piece of paper. Yet it is not identical to a face to face meeting. People tend to concentrate intently on the group work surface (they cannot see each other), not only for tele-data but for the limited sense of tele-presence provided via gestures. People focus attention to objects in the display by pointing at them or by circling objects with the cursor. Drawing and listing are both independent (one person responsible for a drawing) and cooperative (multiple people working together on a drawing). People can—and often do—work simultaneously on any part of the display, and anyone can be actively gesturing, creating, or editing the drawing artifact. The specific observations we had made about GROUPSKETCH are summarized in Table 2. We feel it fair to say that GROUPSKETCH, in spite of its minimalist functionality, is reasonably effective as a distributed work surface.

—Table 2 around here—

2.4: The GroupDraw interface

While GROUPSKETCH is a simple paint program where users can only make and erase marks on a bit-map surface, GROUPDRAW is an object-oriented drawing program with features approaching those of most structured drawing and drafting packages (such as Claris MacDraw). GROUPDRAW was developed as a throw-away experimental platform, and the implementation is buggy and quite fragile. Still, it helped us study and select interface and architecture features for GROUPKIT, our toolkit for building groupware.

¹While XGROUPSKETCH use is similar to GROUPSKETCH, some problems accompany its added features. These will be discussed later in the GROUPDRAW section.

GROUPDRAW supports the multiple active cursors and simultaneous activity that we have seen in GROUPSKETCH. However, users can now create, move, resize, and delete drawing objects (rectangles, lines, circles, text, etc). Also (and similar to XGROUPSKETCH), WYSIWIS has been relaxed to provide a scrollable and resizable work surface. This is a concession to the limited screen real estate available, and also provides users with room to work on private drawings. Figure 3 illustrates a GROUPDRAW screen with three users working on a design. The smaller registration window lists conference participants, their physical location, and their phone numbers.

—Figure 3 around here—

The effect of having objects that can be selected and modified by different people raises a surprising number of interface design issues. The most obvious is what to do when several people try to manipulate an object. The simplest strategy is to let only one person at a time acquire the object—others are prevented from accessing it. Even so, feedback must be supplied to the user differentiating the act of selection from actually acquiring the object, for there could be a time lag between these two states. In GROUPDRAW, users can actively grab an object and start manipulating it—the system assumes optimistically that there will be no conflict. If an acquisition conflict does occur and permission is denied to manipulate the object, the object will snap back to its original position. In practice, system response time is rapid enough that the selection/acquisition process occurs almost instantly—a rejected acquisition appears as a momentary flicker. Still, other visual cues may be necessary to indicate acquisition status during this transition period, such as grey-scale coloring of the object’s handles¹.

Another issue relates to Tang’s fourth design criteria (Table 1) recommending a seamless intermixing of workspace actions and functions. In GROUPSKETCH this was accomplished by designing a relatively modeless interface. The fact that GROUPDRAW users must now select from a variety of object types and go into a particular drawing mode makes this recommendation difficult to fulfil. While the time taken to switch modes is not critical in single-user systems, selecting an object from a palette or menu does detract from the fluidity of a group interaction (this problem is less serious in XGROUPSKETCH because mode changes are infrequent). We do not yet have an acceptable solution, and consider this an area of further research.

¹GROUPDRAW actually allows multiple access to a single object. For example, two people can grab different endpoints of a line and move it. The locking described above is done on the handle level eg when two people try to grab the same endpoint of a line.

A further concern arises from private drawings, which Tang had observed as a positive resource. These drawings are often worked on and then presented to the group at a later time. GROUPDRAW implements privacy in two ways. The first is by providing a scrollable drawing surface; a user can scroll and work in their own area of the screen, and then move the image to the main view (a split screen may work well here, but see below). Second, an object's "coupling status" can be specified (Dewan and Choudhary 1991). Here, a user can indicate whether an object can be manipulated by all others ("all of us can see and touch it"), can only be viewed by others ("I can touch it but you cannot"), or can be private ("only I can see it"). While attaching coupling status to drawing artifacts is a nice technical feature, its actual value (and usability) to real group drawing situations requires further exploration.

Scrolling deserves special attention, for it can completely disrupt the WYSIWIS principle. In GROUPDRAW and in XGROUPSKETCH, views of different participants can now coincide, overlap partially, or be completely disjoint. The result is that a participant is no longer certain of what others can see, and comments such as "Can you see this part of the drawing?" and "Where are you?" occur in the verbal dialog. Strategies exist to mitigate this problem. The first is to supply each user with a separate "radar" window that shows a miniaturized animated view of the complete drawing space and the location of all participants' viewports on it (Smith, O'Shea, O'Malley, Scanlon and Taylor 1989). The cost is the extra screen real estate occupied by the radar window and the cognitive distraction of having to refer to it. A second strategy is *view-slaving*, where the view-port of one participant (called a follower) can be slaved to the view-port of another participant (called a leader) (Pendergast and Hayne 1991). Whenever a leader scrolls the view, the corresponding views of the followers are scrolled as well. What makes view-slaving interesting is that all users can choose to be leaders and followers at the same time, so that anyone can scroll and WYSIWIS is maintained. Of course, "scroll wars" is possible, where people can fight over what part of the image the group should see. However, we expect social protocol will minimize these occurrences.

In summary, our early experience with the GROUPDRAW interface raises more questions than it answers. We know for certain that some of the interaction techniques now found in conventional drawing programming will not transfer well to a multi-user domain.

2.5: Other human factors concerns

Three other human factors concerns are critical to the design of group drawing systems. These are: the need for a better sense of tele-presence; the ability to have a seamless melding of shared work on the computer and the desktop; and consideration of the size of the group.

Does the tele-presence supplied by GROUPSKETCH and GROUPDRAW'S shared cursors, activity space, and voice channel suffice for group work ? As noted in Table 2, GROUPSKETCH users did have a strong focus of attention on the work surface, and we feel that telepresence was adequately supported as long as people were directly interested in the drawing surface. Yet what about the times when the group work surface was not central to the discussion? We conducted several GROUPSKETCH sessions where all users were co-located in a terminal room. When the conversation shifted to a meta-level topic not requiring the drawing surface, we noticed that many of the participants turned away from the computer screen to look at each other instead, even though the layout of this particular room made eye contact awkward. Clearly, people preferred face to face contact in these situations. Smith, O'Shea et al (1989) made a similar observation in a slightly different groupware environment, where two geographically distributed people each had two monitors. One monitor showed the computerized work surface, the other provided a "video tunnel"—a real-time video link that allows eye contact. When the conversation shifted to meta-level discourse, the participants would turn away from the workstation and towards the video monitor. Several other researchers are now examining what happens when the graphical work surface and the video image of the participant's faces are merged via transparent overlays onto a single display (eg ClearBoard-1, Ishii and Arita 1991; Rococco, Scrivner, Clark and Keen 1991). They argue that these systems promote not only eye contact between participants, but also (and perhaps more importantly) "gaze awareness", where a participant can easily recognize what their partner is looking at on the work surface (Ishii and Kobayashi 1992).

A second human factors concern is the barrier that exists between the work we do on our computers and the work we do on our physical desks. We have not yet reached the age where all work is computerized, and many physical items that we wish to share with our geographically-distributed colleagues cannot be easily converted into computer medium in a fast-paced meeting. One simple approach towards a seamless work environment is to allow participants to import static images of documents and other artifacts via paper scanner and video camera/frame grabber; the images can then be manipulated by the groupware [IIS Technologies, date unknown]. Another approach fuses two or more live-video images of computer screens and physical desktop surfaces (Ishii 1990). Here, participants can combine computer artifacts (eg a single-user drawing program) with desktop artifacts (eg an existing drawing). Participants can gesture around and interact with these artifacts using their hands and pencils (by placing them under the video camera and drawing on the surface), or through the standard computer cursor.

Finally, the size of the group will effect how the group uses a drawing surface. The studies of work surface activities mentioned in this paper examined group interactions of two to four people. Yet many face to face meetings involve far larger groups, and much of the technology in the group decision

support system domain accommodates ten to forty or more people (eg GROUPSYSTEMS, Nunamaker, Dennis, Valacich, Vogel and George 1991). Will group drawing tools such as GROUPSKETCH scale up? In one two-hour GROUPSKETCH session, we had eight people working on a design task, a brainstorming task, and an ideas organization task. While we do not have any hard data on the session, the general feeling was that the tool remained useful. However, several issues were noticed. First, eight large cursors cluttered the screen. As a result, we modified GROUPSKETCH so that the cursor size would shrink as the number of group members increased, although any explicit activity (such as drawing or erasing), would restore the active cursor to its full size for the duration of the action. With eight people, the smaller cursors¹ sufficed to provide an uncluttered screen. Participants maintained a group sense of awareness, and temporarily restoring busy cursors to their full size focused the group's attention to the active work. A second issue concerning group size is real estate—large groups need far more working space in view than contemporary bit-mapped displays will provide. Finally, larger groups appeared to spend more time on meta-level discussion. As mentioned previously, this means that the work surface alone will not suffice, and that a visual channel will probably be required as well.

3: Implementation experiences

GROUPSKETCH, XGROUPSKETCH, and GROUPKIT are implemented on Sun workstations running Unix and networked via Ethernet. GROUPDRAW is built upon the Macintosh/AppleTalk platform. The design of all systems contain two features unusual in single-user interface design. First, they are distributed programs—resulting issues are the tradeoffs between a replicated or centralized architecture; the network communication demands; and how users can dynamically register with an existing electronic meeting. Second, all systems support multiple cursors and simultaneous activity—issues here are how multiple cursors are implemented, and how the graphics primitives are structured to support multiple synchronized access.

The internal architecture of GROUPSKETCH and GROUPDRAW are briefly described in this section, indicating how the issues listed above were addressed. XGROUPSKETCH and GROUPKIT are introduced when necessary.

3.1: Replicated vs Centralized Architecture

There are three architectural alternatives for constructing distributed groupware (Figure 4). The *centralized* architecture contains a single central program that controls the distributed work of all users, while a *replicated* architecture executes a copy of the program at every workstation. A *hybrid* approach

¹Cursor sizes here were 8x8 pixels. With larger groups, cursors could be shrunk to occupy only a single pixel on the screen.

combines features of both. While these architectures have been contrasted previously in the design of “collaboration-transparent” view-sharing systems (Ahuja, Ensor and Lucco 1990; Greenberg 1990; Lauwers, Joseph, Lantz and Romanow 1990), many of the technical factors raised are similar to those we have seen in the collaboration-aware groupware discussed in this paper.

—*Figure 4 around here*—

In the centralized approach, a single program, residing on one machine, controls all input and output to the distributed work surfaces (Figure 4a). Server processes residing on each person’s workstation are responsible only for passing user input events to the central program (eg mouse movements, key presses, window resizing), and for displaying output sent to it from the central program. The central program processes these events and then notifies all workstations of display updates. In effect, each workstation is nothing more than a graphical terminal/window server (eg the X window server, Scheifler and Getties 1986), while the central program is one large application program that decides what to do with all users’ input and output. The advantage of a centralized scheme is that synchronization is easy, as state information is consistent since it is all located in one place. The disadvantages are that the complete system is now vulnerable to the failure (either machine or network) of the central agent, and that the central agent could become a network and processor bottleneck as all activity must be channelled through it.

In the replicated approach, there is no central program. Instead, each application program is replicated on every machine (Figure 4b), and the replicated programs stay synchronized by communicating directly with each other. Each replicated program is totally responsible for dealing with its local user’s events, exchanging information with its counterparts on the network, and maintaining the integrity of the drawing surface. The advantages are that network and processor traffic is reduced because communication does not go through a central intermediary, and that the system is more robust to network and machine failure. The catch is that it becomes more difficult to keep the work surfaces and user requests synchronized.

Hybrid approaches are also possible. For example, the participant processes may use a central agent only for synchronization and for mediating potentially conflicting user requests. All other activities are performed within and between participant processes.

An example of a centralized architecture is WSCRAWL 1.0, a public domain group sketching program developed by Brian Wilson at Hewlett-Packard that runs within the X window system. The single

WSCRAWL process collects and decides what to do with all users events (which it collects from the X window server) and where to display the output. Thus it has all the advantages and disadvantages of centralized architectures noted above. Another disadvantage arises when the system follows relaxed WYSIWIS. For example, several user actions in WSCRAWL 1.0, such as menu and control panel selections, are private actions that are visible only on that user's display. The central program must deal with these selections as special cases. Another example is independent scrolling (not implemented in WSCRAWL 1.0). If each user can scroll to different portions of the group drawing surface, the central manager cannot simply broadcast the same output directives to all participants. Rather, it must keep track of each user's view and decide what portion of that view (if any) requires updating. While the programming problems introduced by relaxed WYSIWIS are surmountable, it does add complexity to the code.

In contrast, both GROUPSKETCH and GROUPDRAW use fully replicated architectures, where every workstation has its own local copy of the program. Taking GROUPSKETCH as an example, the replicated GROUPSKETCH processes communicate via Unix stream sockets, using only the eight primitive events listed in Table 3. On startup, the new GROUPSKETCH process asks the registrar daemon (described in the next section) for the location of any other GROUPSKETCH processes. If any exist, the closest process will send the new one a copy of the image, received as an "image transfer" event. This is the only time screen output is ever sent over the communication channel. When any other local user event occurs, such as a cursor movement or line-drawing action, the process will update its own display and pass the event on to the other GROUPSKETCH processes (Table 3). When a remote event is received, it is interpreted and handled in essentially the same manner as a local event, eg remote cursors are moved, remote lines are drawn. Since the only actions that can be done by a GROUPSKETCH process are to either draw/erase on a bit mapped surface or to move the cursor, there is no need to synchronize user activities¹.

—Table 3 around here—

The interaction between replicated processes in GROUPDRAW is more complex. As we have seen, one of the advantages of an object-based drawing system such as GROUPDRAW is the ability to not merely view, but to manipulate the entities in the shared workspace. However, we must ensure that the object's behavior is managed consistently between users, so that (for instance) two people grabbing and dragging the same point on a line do not both succeed. This poses concurrency problems that are not encountered

¹This is not quite true, for it is possible to get out of step. For example, if a user draws on a surface that is simultaneously being erased by someone else, the final appearance of the bitmap could look different depending upon the order in which these events arrive at each participant process. In practise, this is not a problem due to the scarcity of this occurring, and the minimal visual disruption to the drawing.

in the simpler GROUPSKETCH situation. Central architectures can easily resolve this problem as all state information is in one place. However, there are effective approaches to dealing with this under a replicated architecture as well, which we employed in GROUPDRAW. We define an *owner* for each object drawn, who has final authority on all operations affecting the object. The owner of a single object is one of the replicated participant processes in the GROUPDRAW conference¹, and each process may be the owner of more than one object. The responsibility of the owner is to maintain object consistency across all sessions. This is best illustrated by an example.

Suppose three users userA, userB and userC are part of a conference, each running processA, processB, and processC. Within the conference's workspace there is a line object owned by processA. UserB and userC simultaneously select the same end point of the line object—this is a situation where object consistency could be compromised as only one person should be allowed to select the object. Both processB and processC then send a message to processA requesting control of the end point. ProcessA then assigns permission on a first come, first served basis. If processB's request arrives first, processA will send a message to processB granting permission to grab the point, while processC will receive a message denying permission. The owner is then notified when the user has released the point, and the point again becomes available for selection.

Under this approach, object ownership is distributed across all participant processes; it is only the GROUPDRAW session as a whole that maintains full state information for its objects. The result is a fairly robust system. If a participant process leaves (eg when the participant leaves or their node fails), its objects are systematically transferred to other participant processes. When the last person leaves the conference, ownership need not be retained—it is relevant only during a single conference.

A variety of hybrid approaches are possible. In XGROUPSKETCH, for example, we have minimized the need for replicated processes to track each other by moving the responsibility for registration, communication management, and event broadcasting into a centralized manager. As in GROUPSKETCH, the replicated portion of XGROUPSKETCH interprets all local and remote actions, and updates the display accordingly. The difference is that the replicated portion maintains only a single communication channel to the manager, which is used to send local events and to receive remote events. The manager acts as a communication hub, and is responsible for creating and maintaining links to all processes, and for broadcasting any events received.

¹Initially, the owner will be the process supporting the participant who creates the object, whether by drawing it, restoring it from a file, or copying it from another source. The ownership may change during the course of a conference or between conferences by various means.

In GROUPKIT, we are considering (but have not yet implemented) a hybrid approach to the distributed ownership scheme described in GROUPDRAW. A central process (provided as a toolkit component) will mediate concurrency, while the replicated processes will handle all other actions. When a user performs a possibly contentious action, such as grabbing a line's endpoint, the replicated process will ask the central process for permission. If granted, the central process will disallow other users to access that point until it has been informed that the point has been released.

There is no real answer to which scheme works best. Rather, it is a set of tradeoffs that revolve around programming complexity, synchronization requirements, processor speed, the number of participants expected, communication capacity and cost, and so on.

3.2: Registration

We believe that any participant should be able to join and leave the conference at any time. Yet how do people "register" with the shared drawing session, and how is this managed internally? How does each participant process know about and adjust to the comings and going of other participants?

In GROUPSKETCH, a central *registrar* daemon performs dynamic registration functions. The following example indicates how the registrar incorporates new participants into a GROUPSKETCH session.

- 1 An incoming participant or late-comer starts a GROUPSKETCH process, which connects to the registrar, opens a communication port ready for other connections, and sends the port address along with the participant's name and caricature data to the registrar (see Table 3 for the protocol).
- 2 The registrar acknowledges the newcomer and informs all other participants' GROUPSKETCH processes (if any) of the newcomer's address in the network.
- 3 All other GROUPSKETCH processes (if any) connect to the newcomer, with the nearest sending it the current state of the work surface image as an "image transfer" event. The registrar is now out of the loop.

In contrast, GROUPDRAW uses a distributed registration scheme. Each workstation's registrar maintains an AppleTalk socket listener. Whenever a new GROUPDRAW participant enters a session, it announces its arrival over the network. The announcement is heard by every registrar, and connections are made directly. The new entrant then asks one of the other participant processes to send it a display list of the current work surface.

Both schemes work reasonably well. GROUPDRAW'S method has the advantage that it is not tied to a single central registrar. But what is most important in both schemes is that the registrar is fairly independent of the underlying application. It is a high-level toolkit component that should be reusable in other groupware applications requiring registration. In fact, for GROUPSKETCH we were able to reuse in part a registrar that we had originally designed for a completely different application (Greenberg 1990 & 1991). We have included a registrar as a fundamental toolkit component in GROUPKIT.

3.3: Multiple Cursors

Multiple cursor—used for gesturing and tele-presence—present a significant problem, for current window systems only support single cursors. As a result, multiple cursors must be implemented independently from the system-supplied cursors.

In GROUPSKETCH we eschewed window systems completely in favour of a graphical library allowing us to manipulate the bit-map display directly (we used Sun Microsystem's Pixrect library). Large multiple cursors are implemented directly by exclusively OR'ing bitmaps. The general algorithm for handling multiple cursors follows. After initial variable setup and participant registration, the participant process executes a main loop that acts on events arriving from four different sources: the keyboard, the mouse, the registrar, and from other participant processes (via Unix sockets). A participant's local activity is signalled by keyboard events for character insertion; and by mouse events for cursor movement, drawing, and erasing. These events are then broadcast to the remote processes (Table 3); their workspaces are then updated.

As an example of managing cursor movements, assume two GROUPSKETCH participants: userA and userB. When userA moves the cursor, the change is updated immediately on the local screen, and the cursor moved event along with the participant id and the coordinates of the new position are broadcast to userB. When UserB's process receives the event, it looks up the old location of userA's cursor, erases it via an XOR operation, and then redraws it in its new position. The internal table is then updated.

Reading directly from the mouse device driver and writing to the screen provided efficient and fast cursor performance in GROUPSKETCH. However, this approach is costly in terms of implementation effort, for we had to design a graphical interface from the ground-up (eg to manage cursors, menus, simple drawing). We also made a design error in our choice of handling local events differently from remote events. For example, GROUPSKETCH initially had a bug that would cause it to leave a spurious mark on the screen if one cursor is erasing on top of another person's cursor. While we understood why the

problem occurred, the bug took much time to remove due to the different ways local and remote cursor events were handled.

In contrast, GROUPDRAW'S multiple cursors are built in part on top of the Macintosh interface toolkit. While we use standard Macintosh events to determine where the mouse cursor is, we do not use the toolkit routines to display the cursor. Instead, the standard cursor is made invisible, and an XOR scheme similar to the one mentioned above is used in its place. The other difference is that all local and remote events are handled in exactly the same way—the same cursor redrawing function is notified when any cursor is moved (whether local or remote), and the display is updated accordingly. While the standard cursor could have been used to display the local cursor position, we believe it not worth the effort. First, the local cursor has to be treated as a special case from the remote cursors, leading to more complex coding and debugging problems. Second, there are often limitations and system dependencies in application cursors that may conflict with design goals. For example, our cursors are larger than the cursors provided by some window systems. Finally, we have found that our scheme worked well in practice—tracking is fairly smooth and occurs in real time. The same approach was successfully employed in XGROUPSKETCH.

GROUPKIT uses an entirely different approach by providing multiple cursors as transparent overlays to any underlying (programmer-defined) application (Roseman and Greenberg 1992a & 1992b).

GROUPKIT is built on top of the InterViews toolkit (Linton, Calder, Interrante, Tang and Vlissides 1991), and its overlay facility relies upon the InterViews *page* glyph. A page contains one background glyph, and one or more transparent foreground glyphs, each which may be placed independently over the background glyph. The background glyph here represents some application specific graphical component (such as a sketchpad), while the foreground glyphs contain bitmaps representing cursors. As the system cursor is tracked, the foreground glyph for the local cursor is moved, and messages are sent to the other users via the standard mechanisms. The application programmer using GROUPKIT gets multiple cursors for free just by including the specialized cursor page glyph. Similar to cursor management in conventional window systems, the glyph handles the basic display and users' interactions with their multiple cursors, and then passes on cursor movement events to the application. These events can be ignored (which still gives users the ability to gesture around the display) or used by the application for special purposes.

3.4: Communications

People designing real time groupware systems are concerned that the communication channel will be the primary system bottleneck, and often go to great lengths to minimize the information transmitted over the

channel. We believe it has been used as a reason *not* to implement multiple cursors, and for choosing some architectural styles over others (eg Ahuja, Ensor et al 1990; Lauwers, Joseph et al 1990; Lauwers and Lantz 1990).

Our experiences with GROUPSKETCH, XGROUPSKETCH and GROUPDRAW show that communication throughput is not a problem over typical local area networks (Ethernet runs at 10Mbps, Appletalk at ~250Kbps). Our communication requirements were modest. All our implementations use a standard stream-oriented connection that guarantees in-order, error free delivery (via Unix Stream Sockets or Appletalk Data Stream Protocol). When we ran profiling tools on our systems, we were surprised to find that it was the processor speed that was the performance bottleneck. While the network comfortably accommodated the events we had sent over the network, the slower processors had difficulty interpreting all of them in real time.

To get a feel for the network traffic, we traced the number of packets sent by three GROUPSKETCH participants on several short but active design sessions. On average, a total of 15–25 packets/second were transmitted, with an average packet length (ignoring packet overhead) of 5–20 bytes long. This gives a network utilization ranging from 600 to 4000 bits per second. This rate is easily handled by LANs, but could be demanding for a low-speed telephone link. About 75% of the packets sent indicated a “cursor moved” event. Another thing we noted is that each user’s network demands are unequal, for active users generated many more events than inactive users.

Because in our situation CPU processing and not network bandwidth seemed to be the limiting constraint, we adopted the following philosophy. Each participant process sends out as much information as possible to the other processes. For example, cursor movement is broadcast as often as possible. On the receiving end, decisions must be made about what information to process when backlogs in the event queue occur. A slower machine, for example, might ignore all cursor updates received from a particular process except for the last. While this can result in jerky cursor motion, it is far less disconcerting than waiting for the cursor to “catch up”.

Of course, there will still be situations where the network is the bottleneck. Delays due to latency, for example can seriously undermine performance. Also, network software may not be designed to handle the frequent short messages used to broadcast cursor events. For example, Microsoft Windows has an induced delay in packet transmission, probably due to the way its kernel only processes interrupts after a certain arbitrary number of clock ticks (Hayne, Pendergast and Greenberg 1992). For pragmatic reasons, transmitting all cursor events may not be feasible, and other cursor transmission and display

strategies may be required. Hayne, Pendergast and Greenberg (1992) offer four low-bandwidth alternatives to full motion cursors.

Point – when the user points and presses a mouse button, a single message is sent to all remote users. The gesturing cursor will jump to the new location.

Point and quiver – in order to attract the viewer’s attention, the above strategy is modified so that the receiving station *jitters* the cursor back and forth in its new position.

Limited motion – Only a certain number of gesture messages per second are allowed between stations; the fewer messages sent, the jerkier the motion (e.g. the SHDR group sketching system, by Paul Dourish, Rank Xerox EuroParc).

Limited motion with smoothing – After receiving a new cursor position from the sender, the receiving station computes several points in between the previous and current location, and then animates the cursor through these intermediate points.

One final point worth mentioning is the relation of bandwidth to n , the number of participants. In the current replicated architecture, the worst case network demands is a function of $n*(n-1)$, which occurs only when every participant is active on the work surface. In this case, each replicated process must broadcast its local events to all others (Figure 4b). Of course, simultaneous activity from all members of larger groups is unlikely, and the average performance demands will be much smaller. It should be noted that if a multi-cast network were available, then the communication demands would be linear to the number of participants, for only a single message need be broadcast by the sender.

3.5: Drawing primitives

A structured drawing package provides its users with a set of drawing primitives: lines, circles, rectangles and so on. While these are familiar to most simple graphics packages, the property that they are sharable is still fairly novel. This section shows how object-oriented programming is used in GROUPDRAW to isolate most multi-user characteristics into a prototypical drawing object. Sub-classes, which inherit these characteristics, need only specify the actual graphical properties of the drawable object. This prototypical object will be discussed in terms of its properties and operations. We will indicate how its subclasses can be implemented.

—Table 4 around here—

All drawable objects have certain properties which we isolate as much as possible into a root object called *GroupGraphicalObject*. Table 3 provides some detail of the instance variables and methods it contains. Each graphical object has an *ownerProcess*, which is by default the creator of the object. The *ownerProcess* serves to arbitrate contention in manipulating the object. Objects also have a

couplingStatus (Dewan and Choudhary 1991) that indicates the extent to which graphical objects are shared. As mentioned in Section 2.4, GROUPDRAW defines three coupling levels: private, public, and sharable. Each object is also referenced by a unique *id*, which is used in all network messages to identify the object being manipulated.

To create a new object, whether it be in response to a local or remote drawing action, we provide a standard initialization method *initalizeGroupObject()*, which will set the instance variables mentioned above. Subclasses will specialize this method to initialize any extra properties it may have eg zeroing out the endpoints of a line (see *LineObject*, Table 3). For communication and storage purposes, each object must be able to construct and interpret a string representing itself (*make/sendDescription()*); this is used to save and restore images and to send update information to new users in the conference. The latter is implemented by requesting each object to tell the new user about itself.

Next, we look at operations dealing with changing the coupling status of an object. GROUPDRAW insists that the owner approve status changes; processes request permission by the *requestChangeStatus()* method. If the owner grants permission, the *doChangeStatus()* method will actually change the object status and broadcast the change to all participants. Table 3 lists several other methods that follow this request/do form of arbitration.

Because multiple users may select and start to drag an object asynchronously, the object may be in slightly different places on different screens. Yet each selected object will want to tell the other processes where it had been selected. Passing the precise pixel coordinates is often meaningless, since that point may not match its partner on the remote object. Instead we define a “logical” point for each object. For example, the logical points of a line will be the two endpoints (*start/endPoint*). Dragging a point within an object would then be described in relation to these logical points (*whereGrabbed*). Most of the group interaction algorithms are handled within the *GroupGraphicalObject*. The root methods handle arbitration for object selection (*requestToGrab()*, *doGrab()*, *drag()*, *endGrab()*) The *doDrag()* method, specialized for each subclass, actually does the dragging— it does not need to know how other users are manipulating the object. The actual drawing of the object by *draw()* will, of course, be specialized to each sub-class.

What must a sub-classed object such as *LineObject* be responsible for? First, it needs to create or interpret the description string it defines if it is to send a complete object description over the network. Second, given a physical point, it must determine the corresponding logical point. Third, it must handle the object-specific graphical activities, including methods to draw and erase itself, to drag itself around,

and to save enough state information to undo the dragging operation if the object's owner refuses dragging permission (*save/restoreOriginal()*). Object status, ownership, contention, and other issues need not be dealt with by the sub-class.

4: Conclusions

This paper introduced both human and technical factors we have experienced designing four multi-user systems: the GROUPSKETCH and XGROUPSKETCH sketchpads, the GROUPDRAW object-oriented drawing package, and the GROUPKIT groupware toolkit.

On the human factors side, it may appear that some of the design principles mentioned in Section 2 are self-evident eg multiple cursors for gesturing, allowing simultaneous activities, and so on. Yet there are many examples of related groupware systems that have failed to live up to these seemingly self-evident criteria. Consider Xerox PARC's BOARDNOTER, a computerized whiteboard used to support face to face meetings (Stefik, Bobrow, Foster et al 1987; Stefik, Foster, Bobrow et al 1987). While a single large tele-pointer could be seen by all, individual cursors were not. Neither did participants see each others actions as they occurred, for actions were not broadcast until a complete graphical stroke was made or a complete text line entered. XSKETCH, a recent object-based group drawing package suffers a similar lack as its objects are only transmitted after they are created (Lee 1990). WSCRAWL 1.0, mentioned in Section 3.1, does not show multiple cursors¹. Group Technologies' ASPECTS does not necessarily show multiple cursors, nor can the process of creating or manipulating an object be seen by participants. We have also seen several other systems now under development that fail to meet the basic human factors requirements of a group drawing tool.

On the positive side, there are several systems (including X/GROUPSKETCH and GROUPDRAW) that do support the kinds of interactions people expect from a group drawing surface. All have one thing in common: they were derived from Tang's design principles as listed in Table 1. While these systems are quite diverse, they all share a common feel, and observations of use are strikingly similar. Three systems, for example, are video based and work by fusing video images: VIDEODRAW (Tang and Minneman 1990), VideoWhiteboard (Tang and Minneman 1991) and TEAMWORKSTATION (Ishii 1990). All are limited by scalability, for serious image deterioration results when too many video images are fused. In contrast, COMMUNE is a workstation-based multi-user sketchpad built independently but in parallel with GROUPSKETCH (Bly and Minneman 1990; Minneman and Bly 1990 & 1991). Although the interface to the two systems are remarkably similar, there are some minor differences. In COMMUNE,

¹WSCRAWL 2.0 has just been released. It is a completely new design and implementation, and now how multiple cursors and excellent functionality. The creator cites the influence of Bly and Tang on their new work.

people use a stylus to write directly on top of the horizontally-oriented monitor—the resulting artifacts are superior to the ones generated on our mouse-based system. Other workstation-based sketchpads are now coming available, for example, CAVEDRAW (Lu and Mantei 1991); WSCRAWL 2.0, SHDR from Rank EuroPark, ROCOCO SKETCHPAD (Scrivner, Clark et al 1991); MUGE (Pendergast and Hayne 1991); GROUPTELECONFERENCING SYSTEM (IIS Technologies, date unknown). Some include video-links to promote face to face contact. We also believe that these guidelines are generalizable beyond group drawing. For example, we know of one group brainstorming system (Cognoter) that has been redesigned to fit these criteria (Tatar, Foster and Bobrow 1991).

We have also shared our experiences implementing a computational shared drawing surface, concentrating on where a multi-user drawing application would differ from its single-user counterparts. We found that there are tradeoffs in the choice between replicated, centralized, and hybrid architecture, and that the reason to choose one style of another will often depend upon the physical requirements of the system. We recommended that conference registration be managed as independently as possible from the underlying application, and that it is best handled as a high-level toolkit component. Multiple cursors are considered fundamental to these systems; we recommended that future interface toolkits and window systems support these directly, as we now do in GROUPKIT. We have also found that communication bandwidth on moderate speed local area networks is not a problem. While we recognize that slow-speed telephone lines are still a fact of life, we suggested that in general the underlying system functionality should not be compromised for communications problems that may not exist (especially since 14,400 bps modems are now available). Finally, we outlined how a multi-user graphics library can be created by having most of its collaborative-aware properties reside within a root prototypical graphics object. By sub-classing, it should be fairly straight forward to extend the library via conventional graphics procedures.

Acknowledgement. This research is supported by the National Science and Engineering Research Council of Canada.

Software Availability. GROUPSKETCH is now available via anonymous ftp from the University of Calgary, Department of Computer Science via *cpsc.ucalgary.ca*. Contact Saul Greenberg for further information on the availability of XGROUPSKETCH and GROUPKIT.

References.

- Ahuja, S. R., Ensor, J. R. and Lucco, S. E. (1990) "A comparison of applications sharing mechanisms in real-time desktop conferencing systems." In *Proceedings of the Conference on Office Information Systems*, pp. 238-248, Boston, April 25-27.
- Bly, S. (1988) "A use of drawing surfaces in different collaborative settings." In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '88)*, pp. 250-256, Portland, Oregon, September 26-28, ACM Press.
- Bly, S. A. and Minneman, S. L. (1990) "Commune: A shared drawing surface." In *Proceedings of the Conference on Office Information Systems*, pp. 184-192, Boston, April 25-27.
- Dewan, P. and Choudhary, R. (1991) "Flexible user interface coupling in collaborative systems." In *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 41-48, New Orleans, April 28-May 2, ACM Press.
- Dykstra, E. A. and Carasik, R. P. (1991) "Structure and support in cooperative environments: The Amsterdam Conversation Environment." *Int J Man Machine Studies*, **34**(3), pp. 419-434, March. In the special edition on CSCW & Groupware. Republished in Greenberg, 1991.
- 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 (CSCW '88)*, pp. 13-24, Portland, Oregon, September 26-28, ACM Press.
- Gaines, B.R. (1991) "Modeling and forecasting the information sciences." *Information Sciences*, 57-58, pp. 3-22.
- Garner, S. W., Scrivner, S. A. R., Clarke, A. A., Clark, S., Connolly, J. H., Palmen, H., Schappo, A. and Smyth, M. G. (1991) "The use of design activity for research into computer supported cooperative work (CSCW)." In *Proceedings of DATER '91*, pp. 84-96, Loughborough, UK, September.
- Greenberg, S. (1990) "Sharing views and interactions with single-user applications." In *Proceedings of the ACM/IEEE Conference on Office Information Systems*, pp. 227-237, Cambridge, Massachusetts, April 25-27.
- Greenberg, S. (1991) "Personalizable groupware: Accomodating individual roles and group differences." In *Proceedings of the European Conference of Computer Supported Cooperative Work (ECSCW '91)*, pp. 17-32, Amsterdam, September 24-27, Kluwer Academic Press.
- Greenberg, S. and Bohnet, R. (1991) "GroupSketch: A multi-user sketchpad for geographically-distributed small groups." In *Proceedings of Graphics Interface '91*, pp. 207-215, Calgary, Alberta, June 5-7.
- 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.

- Greenberg, S., Roseman, M., Webster, D. and Bohnet, R. (1992) "Issues and experiences designing and implementing two group drawing tools." In *Proceedings of Hawaii International Conference on System Sciences*, 4(, Kuwail, Hawaii, January, IEEE Press.
- Grudin, J. (1989) "Why groupware applications fail: Problems in design and evaluation." *Office: Technology and People*, 4(3), pp. 245-264.
- Hayne, S. C., Pendergast, M. O. and Greenberg, S. (1992) "Gesturing through cursors: Implementing multiple pointers in group support systems." Working Paper, University of Calgary, Department of MIS, July.
- IIS Technologies (date unknown) "Introduction to the Group TeleConferencing System." IIS Technologies Inc.
- Ishii, H. (1990) "TeamWorkStation: Towards a seamless shared space." In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '90)*, pp. 13-26, Los Angeles, California, October 7-10, ACM Press.
- Ishii, H. and Arita, K. (1991) "ClearFace: Translucent multiuser interface for TeamWorkstation." In *Proceedings of the European Conference of Computer Supported Cooperative Work*, Amsterdam, September, Kluwar press.
- Ishii, H. and Kobayashi, M. (1992) "ClearBoard: A seamless medium for shared drawing and conversation with eye contact." In *ACM Conference on Human Factors in Computing Systems*, pp. 525-532, Monterey, California, May 3-7, ACM Press.
- Johansen, R. and Bullen, C. (1984) "Thinking ahead: What to expect from teleconferencing." *Harvard Business Review*, pp. 4-10, March/April. Reprinted in Greif, 1988.
- Lauwers, J. C., Joseph, T. A., Lantz, K. A. and Romanow, A. L. (1990) "Replicated architectures for shared window systems: A critique." In *Proceedings of the Conference on Office Information Systems*, pp. 249-260, Boston, April 25-27.
- Lauwers, J. C. and Lantz, K. A. (1990) "Collaboration awareness in support of collaboration transparency: Requirements for the next generation of shared window systems." In *Proceedings of the ACM/SIGCHI Conference on Human factors in Computing*, Seattle Washington, April 1-5, ACM Press.
- Lee, J. J. (1990) "Xsketch: A multi-user sketching tool for X11." In *Proceedings of the Conference on Office Information Systems*, pp. 169-173, Boston, April 25-27.
- Linton, M.A., Calder, P.R., Interrante, J.A., Tang, S. and Vlissides, J.M. (1991) "InterViews Reference Manual." Stanford University, California, Spetember.
- Lu, I. and Mantei, M. (1991) "Idea management in a shared drawing tool." In *Proceedings of the European Conference of Computer Supported Cooperative Work (ECSCW'91)*, pp. 97-112, Amsterdam, September 24-27, Kluwer Academic Press.

- Minneman, S. L. and Bly, S. A. (1990) "Experiences in the development of a multi-user drawing tool." In *The 3rd Guelph Symposium on Computer Mediated Communication*, pp. 154-167, Guelph, Ontario, Canada, May 15-17, University of Guelph Continuing Education Division.
- Minneman, S. L. and Bly, S. A. (1991) "Managing a trois: A study of a multi-user drawing tool in distributed design work." In *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 217-224, New Orleans, April 28-May 2, ACM Press.
- MIT (1983) "Talking heads." In *Discursions*, Boston, Mass, Architecture Machine Group, MIT. Optical disc.
- Nunamaker, J. F., Dennis, A. R., Valacich, J. S., Vogel, D. R. and George, J. F. (1991) "Electronic meeting systems to support group work." *Comm ACM*, **34**(7), pp. 40-61, July.
- Pendergast, M. O. and Hayne, S. C. (1991) "Assisting groups during the merging process." In *Decision Sciences Conference*, November.
- Roseman, M. and Greenberg, S. (1992a) "GroupKit: A groupware toolkit." In *ACM SIGCHI Conference on Human Factors in Computing Systems, Proceedings of Posters and Short Talks*, pp. 43, Monterey, California, ACM Press. May 3-7.
- Roseman, M. and Greenberg, S. (1992b) "GroupKit: A groupware toolkit for building real-time conferencing applications." In *ACM Conference on Computer Supported Cooperative Work (CSCW '92)*, Toronto, Ontario, November 1-4, ACM Press.
- Sarin, S. and Greif, I. (1985) "Computer-based real-time conferencing systems." *IEEE Computer*, **18**(10), pp. 33-45. Reprinted in Greif, 1988.
- Scheifler, R.W. and Gettys, J. (1986) "The X window system." *ACM Transactions on Graphics*, **5**(2), pp. 79-109, April.
- Scrivner, S. A. R., Clark, S. M. and Keen, N. L. (1991) "The role of workspace-replication in the development of remote CSCW systems." , LUTCHI Research Centre, Dept of Computer Studies, Loughborough University, Loughborough, Leicestershire, UK. Submitted to CHI '92.
- Smith, R. B., O'Shea, T., O'Malley, C., Scanlon, E. and Taylor, J. (1989) "Preliminary experiences with a distributed, multi-media,problem environment." In *Proceedings of the 1st European Conference on Computer Supported Cooperative Work (EC-CSCW '89)*, Gatwick, U.K., September 13-15, Computer Sciences House, Sloug, UK.
- Stefik, M., Bobrow, D. G., Foster, G., Lanning, S. and Tatar, D. (1987) "WYSIWIS revised: Early experiences with multiuser interfaces." *ACM Trans Office Information Systems*, **5**(2), pp. 147-167, April. An earlier version appeared in CSCW '86.
- Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. and Suchman, L. (1987) "Beyond the chalkboard: Computer support for collaboration and problem solving in meetings." *Comm ACM*, **30**(1), pp. 32-47. Reprinted in Greif, 1988.

- Tang, J. C. (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 Palo Alto Research Center, Palo Alto, California.
- Tang, J. C. (1991) "Findings from observational studies of collaborative work." *Int J Man Machine Studies*, **34**(2), pp. 143-160, February. In the special edition on CSCW & Groupware. Republished in Greenberg, 1991.
- Tang, J. C. and Leifer, L. J. (1988) "A framework for understanding the workspace activity of design teams." In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '88)*, pp. 244-249, Portland, Oregon, September 26-28, ACM Press.
- Tang, J. C. and Minneman, S. L. (1990) "Videodraw: A video interface for collaborative drawing." In *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 313-320, Seattle Washington, April 1-5, ACM Press.
- Tang, J. C. and Minneman, S. L. (1991) "VideoWhiteboard: Video shadows to support remote collaboration." In *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 315-322, New Orleans, April 28-May 2, ACM Press.
- Tatar, D. G., Foster, G. and Bobrow, D. G. (1991) "Design for conversation: Lessons from Cognoter." *Int J Man Machine Studies*, **34**(2), pp. 185-210, February. In the special edition on CSCW & Groupware. Republished in Greenberg, 1991.

List of Table Captions

Table 1 Six criteria for designing a communal work surface (condensed from Tang 1989), and how they were met in GROUPSKETCH

Table 2. Observation of GROUPSKETCH use.

Table 3. Communication protocol between GROUPSKETCH processes

Table 4. GROUPDRAW'S GroupGraphicalObject and its LineObject sub-class

List of Figure Captions

Figure 1: A sample GROUPSKETCH session

Figure 2: Artists Jacqui MacFarland and Colleen Campbell produce a joint drawing with XGROUPSKETCH. The system runs within an X window; the upper panel shows controls for font selection, line thickness, color choice, and so on.

Figure 3: A sample GROUPDRAW session, showing the work surface and the registration window

Figure 4: Schematic of a centralized, replicated and hybrid architecture

Design Criteria	Reasons	How criteria were met in GROUPEKETCH
<p>1) Provide ways of conveying and supporting gestural communication.</p> <p>Gestures should be clearly visible, and should maintain their relation with objects within the work surface and with voice communication.</p>	<ul style="list-style-type: none"> • gestures are a prominent action • gestures are typically made in relation to objects on the work surface • gestures must be seen if they are to be useful • gestures are often accompanied by verbal explanation 	<p>Physical gestures are conveyed via specialized multiple cursors synchronized with voice.</p> <ul style="list-style-type: none"> • Since gestures must be seen in order to convey information, all cursors within a work surface are always visible to all participants. Cursors are also made prominent on the display by their larger than normal size. • Cursors change their shape to reflect a natural action. Four gesture modes are supported (pointing, writing, erasing, and directing attention) by distinct cursor shapes (Figure 1). The default cursor shape is the pointing hand, while the large arrow allows users to point at and direct the group's attention with greater emphasis than the normal hand. • Cursors are unique, each identifying the person it belongs to by labelling it with the user's name. In addition (and more subtly) each cursor is orientated at different angles (Figure 1). • Cursor movements appear with no apparent delay on all displays, which means that they remain synchronized with verbal communication. • Cursors always maintain their same relative location on every display so that they retain their relation to the work surface objects.
<p>2) Minimize the overhead encountered when storing information.</p>	<ul style="list-style-type: none"> • only one person usually records information • other participants should not be blocked from continuing private or group work while information is being stored 	<ul style="list-style-type: none"> • Any person may store a snapshot of the current work surface into their own private directories at any time. While that particular workstation display will 'freeze' for a short period, other workstations remain unaffected. Any person may restore their private images back to the public work surface at will.
<p>3) Convey the process of creating artifacts to express ideas.</p>	<ul style="list-style-type: none"> • the process of creation is in itself a gesture that communicates information • speech is closely synchronized with the creation process • artifacts in themselves are often meaningless 	<ul style="list-style-type: none"> • In GROUPEKETCH, any work surface action, no matter how small, is visible with no apparent delay on all participants' screens. Every movement of the cursor, every pixel that is drawn, and every letter typed is immediately broadcast to other screens and are therefore immediately visible. • As with cursors, all details of artifact creation and manipulation are transmitted in real time and remain synchronized with accompanying speech.
<p>4) Allow seamless intermixing of work surface actions and functions</p>	<ul style="list-style-type: none"> • a single action often combines aspects of listing, drawing and gesturing • writing and drawing alternates rapidly • actions often address several functions 	<ul style="list-style-type: none"> • The simplicity of GROUPEKETCH allows it to have a nearly modeless interface. When no mouse buttons are depressed, the cursor is in the pointing gestural state. Drawing occurs as long as the left button is depressed, which also turns the cursor into a pen. Similarly, the right button will invoke the eraser function, and the middle button the large arrow for focussing attention. Typing immediately inserts text at the current cursor location and the cursor image changes to the pen, automatically reverting back to the hand-shaped cursor after a reasonable pause in typing is detected.

Table 1. Six criteria for designing a communal work surface (condensed from Tang 1989), and how they were met in GROUPEKETCH (continued on next page)

... continued from previous page

Design Criteria	Reasons	How criteria were met in <i>GroupSketch</i>
<p>5) Enable all participants to share a common view of the work surface while providing simultaneous access and a sense of close proximity to it.</p>	<ul style="list-style-type: none"> • people do not see the same things when orientation differs • simultaneous activity is prevalent • close proximity to the work surface encourages simultaneous activity 	<ul style="list-style-type: none"> • Advantages of a common view were considered more beneficial than the lesser advantages of differing orientations. By using a “what you see is what I see” (WYSIWIS) display and by having all cursors present, we believe GROUPSKETCH promotes a close sense of proximity. As participants track other cursors, they naturally associate actions in the work surface with people who are executing those actions. • Simultaneity is fully supported. All participants have free and equal access to the work surface. Any one can do anything at any time.
<p>6) Facilitate the participants’ natural abilities to coordinate their collaborations</p>	<ul style="list-style-type: none"> • people are skilled at coordinating communication • we do not understand the coordinating process well enough to mechanize it 	<ul style="list-style-type: none"> • As GROUPSKETCH does not enforce any style of social protocol and as all participants are in direct control of their actions, the group is free to use whatever coordination method suits them (an argument favouring this approach is presented by Dykstra and Carasik 1991).

Table 1 Six criteria for designing a communal work surface (condensed from Tang 1989), and how they were met in GROUPSKETCH

Observation	Details
GROUPSKETCH is very easy to learn.	People with even limited computer experience learnt GROUPSKETCH in moments (less than a minute). We attribute this ease to its direct analogy to the paper sketchpad, the modeless nature of the system, and its simple syntax.
GROUPSKETCH is effective.	In spite of its simplicity, GROUPSKETCH worked. Participants were able to pursue their tasks effectively, using strategies analogous to those observed in face to face design meetings.
The worst part of GROUPSKETCH is trying to draw with a mouse.	People expressed frustration when drawing with a mouse. A stylus would have been a large improvement.
Increasing the number of participants in an open floor policy increases parallel activity but also decreases focused attention.	We observed much simultaneous activity. As noted by Tang (1989), this comes at the price of reduced group attention. For example, when four participants were collaborating, one person commented that she found it difficult to listen to another participant when others were actively writing or drawing in the communal work area. We expect this problem to be exacerbated as group size increases. Yet most participants agreed that restricting access to the work surface or introducing turn-taking would be unacceptable.
Movement of the cursor synchronized with a participant's voice provides the greatest sense of tele-presence.	The presence of even idle cursors in the work surface was considered important by participants. People did not have serious problems distinguishing who was doing what. Still, the quality of presence did not match that of a face to face meeting. For example, we observed two occasions when visually separated but co-located participants involved in an intense discussion left their computers to speak face to face.
The shared work surface captured participants' attention and focused interaction.	There is a strong focus of attention on the work surface. Participants' eyes remained fixed on the shared area for long periods of time, as if they did not want to miss any of the actions occurring in the work surface. The ease of drawing and talking simultaneously around artifacts seemed to provide a focused interaction.
Participants desired greater functionality.	People familiar to computer systems wanted functionality greater than a simple sketchpad could provide. These included object-oriented drawing tools over free-hand bit-mapped sketching, editable text fields, and other features commonly available in single-user graphical packages. This finding was our main motivation for designing <i>GroupDraw</i> .
Intermixing listing and drawing (text and graphics) occurred frequently and naturally.	Resulting artifacts contained a good mixture of graphics and textual lists of points.
Vertical orientation of the work surface removed the physical limitations of the table top.	Users had no problem recognizing objects on the display. As people could literally draw on top of one another, we observed people working together on objects in quite close proximity (examples include multiple people erasing different parts of a single line and cooperative construction of a drawing artifact).
Saving only one image is not enough.	GROUPSKETCH only allowed any one user to save one image at a time. This was not enough to allow rapid switching between drawings.
The work surface is too small.	The work surface quickly becomes cluttered during long design sessions, especially with larger group sizes. Larger displays, windowing strategies, or better storage and retrieval facilities are required.

Table 2. Observation of GROUPSKETCH use.

Event	Information passed
Registering a new user	host name, port number, name of participant, caricature
Unregistering a user	Id of participant
Cursor moved	Id of participant, cursor shape, new coordinates
Drawing a line	Id of participant, start and end coordinates
Erasing a region	Id of participant, coordinates of region
Listing	Id of participant, string location, string, cursor shape, location of cursor
Clearing screen	—
Image transfer	binary data of the work surface image

Table 3. Communication protocol between GROUPSKETCH processes

GroupGraphicalObject, the root object of all graphics primitives

<i>instance variables</i>	
int couplingStatus	-indicates if object is private, public, or shareable
int ownerProcess	-indicates who the owner of the process is
int id	-a unique reference that identifies the object across the whole system
boolean acquired	-indicates if the object is currently being manipulated by a participant
point whereGrabbed	-a logical point indicating the position where the object was acquired
<i>methods for initializing and destroying an object, and for sending its description to others</i>	
initializeGroupObject()	-initializes the object, filling in any defaults
destroyObject()	-destroys the object and frees its space
makeDescription()	-make a description of the object suitable for transmission over the network
sendDescription()	-send the description to a requestor
<i>methods for changing object attributes</i>	
requestChangeStatus()	-request the object owner to change the coupling status
doChangeStatus()	-actually change the status, and broadcast change over network
requestChangeOwner()	-request the object to change its owner
doChangeOwner()	-actually change the owner, and broadcast change over network
<i>methods for graphically manipulating and drawing the object</i>	
draw()	-a place holder for a routine that will draw the object on the screen
requestToGrab()	-request the object for permission to acquire it
doGrab()	-permission is granted or denied
endGrab()	-an acquired object is relinquished
saveOriginal()	-the original position of a selected object
restoreOriginal()	-restore a moved object to its original position
drag()	-high level drag handler; checks permission, broadcasts changes, etc
doDrag()	-a place holder; this routine will actually drag the object
whereGrabbed()	-a place holder; checks to see where object was grabbed

LineObject, a sub-class of GroupGraphicalObject

<i>instance variables</i>	
point startPoint	-the start point of the line
point endPoint	-the end point of the line
<i>methods</i>	
initializeGroupObject()	-initialize line-specific attributes, then call the super-class' initializeGroupObject
makeDescription()	-specialized to include line descriptions, then calls the inherited method
saveOriginal()	-specialization line-aware form of the inherited method
restoreOriginal()	-specialization line-aware form of the inherited method
doDrag()	-actually drags the object
draw()	-actually draw the object

Table 4. GROUPDRAW'S GroupGraphicalObject and its LineObject sub-class

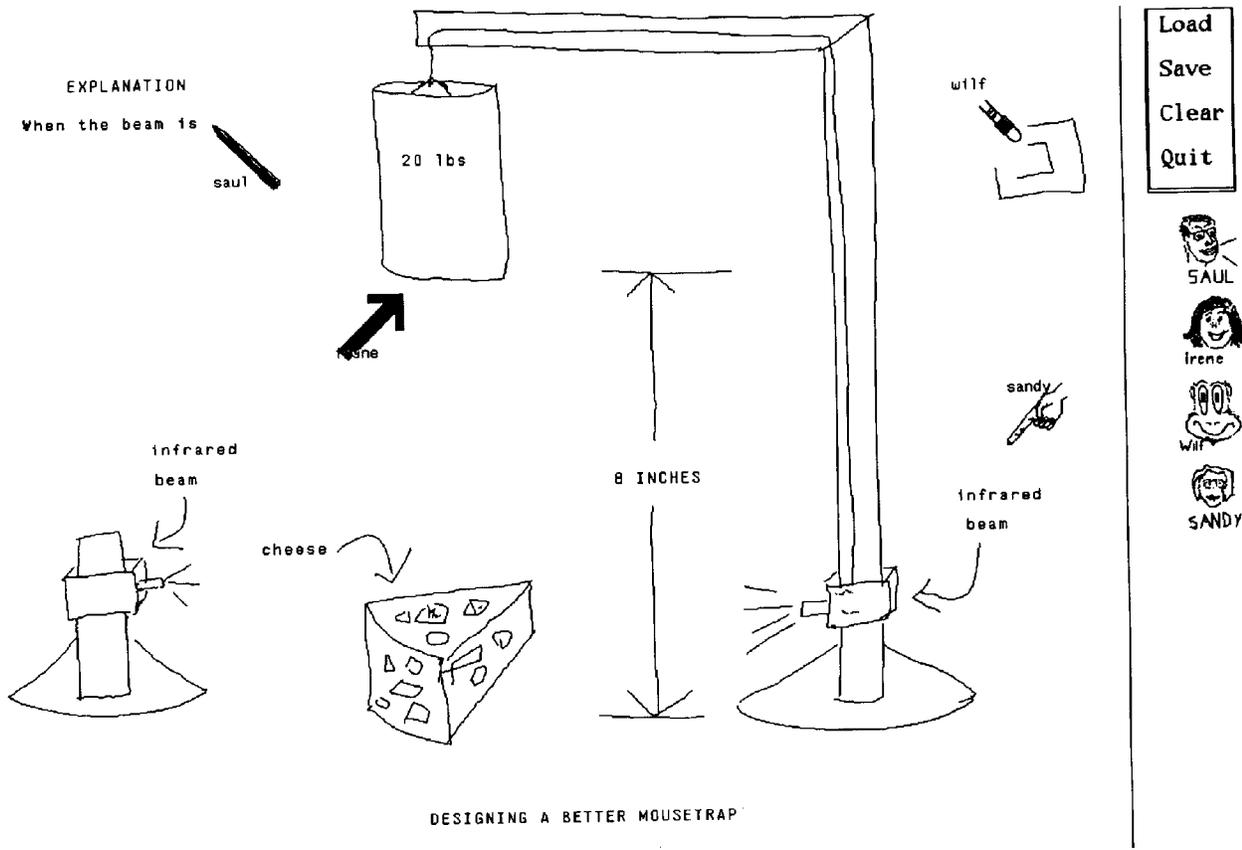


Figure 1: A sample GROUPSKETCH session
(THIS IS A ROUGH COPY OF THE FIGURE IN THE ACTUAL PAPER)

----- FIGURE 2 NOT AVAILABLE IN THIS DOCUMENT ---

Figure 2: Artists Jacqui MacFarland and Colleen Campbell produce a joint drawing with XGROUPSKETCH. The system runs within an X window; the upper panel shows controls for font selection, line thickness, color choice, and so on.

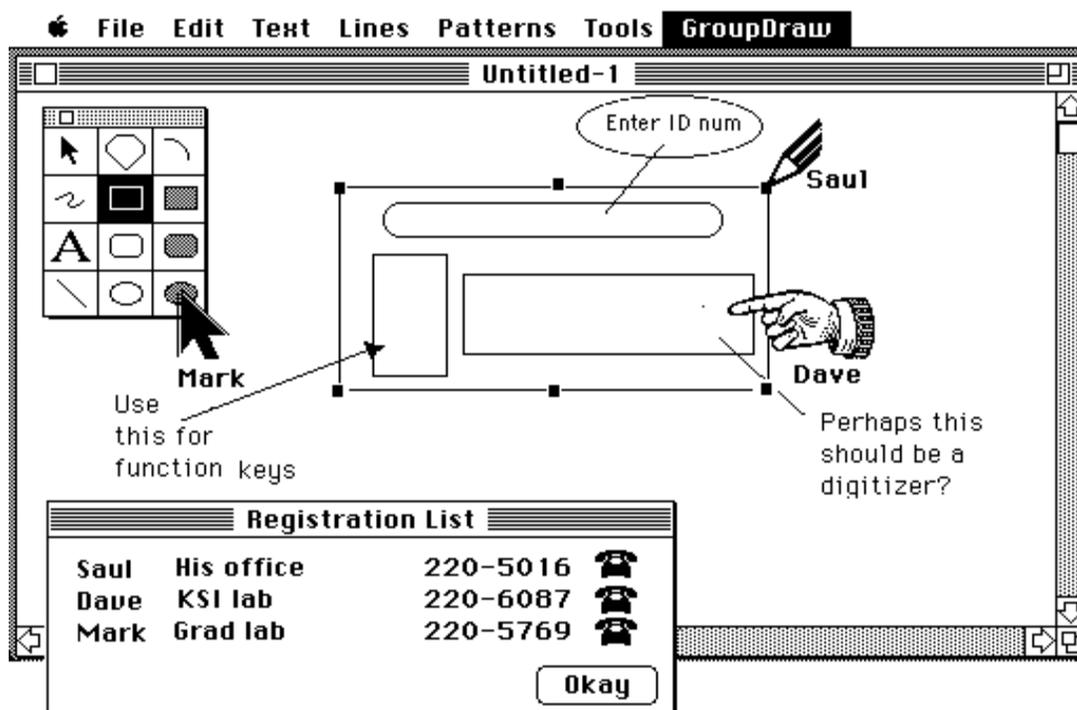


Figure 3: A sample GROUPDRAW session, showing the work surface and the registration window

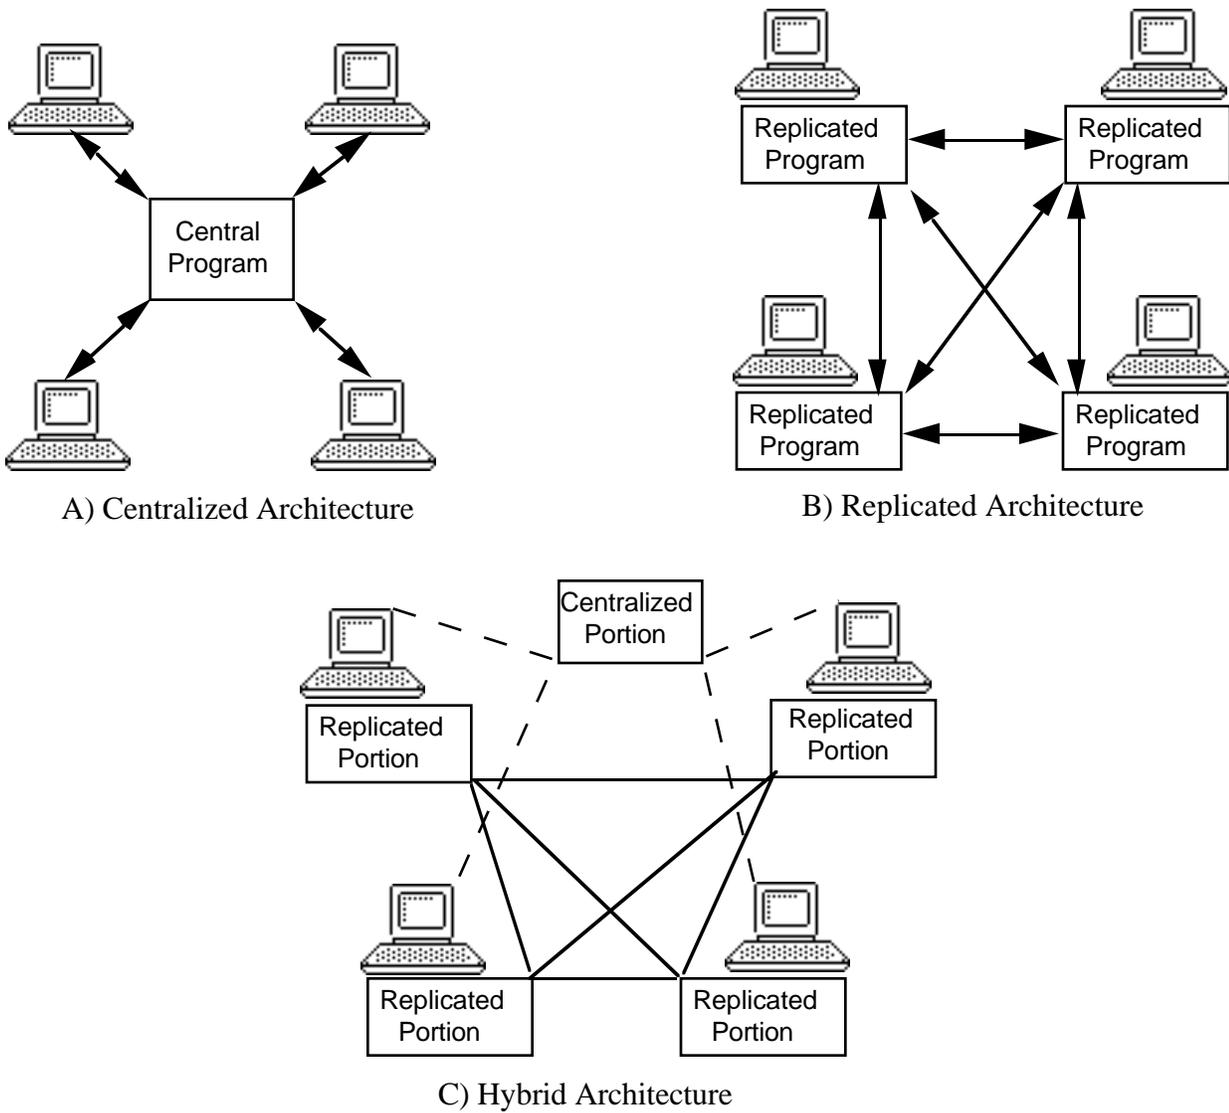


Figure 4: Schematic of a centralized, replicated and hybrid architecture