Managing à trois: a study of a multi-user drawing tool in distributed design work

Scott L. Minneman
Sara A. Bly

Systems Sciences Laboratory
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304

Minneman@parc.xerox.com, (415) 494-4353
Bly@parc.xerox.com, (415) 494-4360

ABSTRACT

A multi-user drawing tool was used by participants in a distributed design exercise conducted in a multi-media working environment. The goal of the study was to explore how observations from our earlier studies of shared drawing in two-person design activity would hold up when three participants worked together. Additionally, the study provided opportunities to contrast video/audio connections with audio-only connections and to discover new behaviors that emerge in the use of new technologies.

Participants successfully used the shared drawing system with no observed difficulties attributable to the addition of a third user. Audio-only connections appeared to adequately support this work activity, but details of the participants' interactions in the exercise raised questions that deserve further study. Finally, observations suggest that drawing tools such as the one reported here may offer support for alternative forms of participation in collaborative work.

KEYWORDS: shared drawing, collaboration, group work, distributed work, video

INTRODUCTION

A variety of recent research and commercial development efforts have been aimed at aspects of supporting shared drawing activity in distributed situations. These efforts have a broad spectrum of foci, from demonstrations of novel means of sharing a distributed workspace to explorations of the programming issues involved in shared applications. We have found studies of the use of the tools that have been developed over the years, or even more general studies of graphical activity in distributed working groups, to be considerably less prevalent.

We developed Commune, a shared drawing surface prototype, as a means of providing support particularly for distributed design groups. Our interests in this area are threefold: developing tools to support distributed work groups generally, studying the process of designing a collaborative work tool based on observations of real work activity, and understanding how collaborators can do design work in distributed settings. Initial Commune studies concentrated on pairs of collaborators involved in short-duration design work. One goal in the research reported here is to begin exploring how some of our observations hold up when the activity is expanded to include additional participants.

In our earlier Commune research the distributed work setting was comprised of the prototype Commune workstations, video connections, and full-duplex audio. One attractive aspect of Commune is that shared drawing capability can be provided without extensive infrastructure support by the utilization of modems over dial-up phone lines. Thus, another goal in the research reported here was to gain experience with collaborators using Commune in a low-bandwidth situation without the video communication channel.

Commune is one of a number of shared drawing tools being prototyped at the Xerox Palo Alto Research Center (PARC). A major research concern in the design and development of these tools is how to ground design decisions in observations of actual or naturalistic work activity [Tatar, 1989]. An important aspect of that work is the iteration through

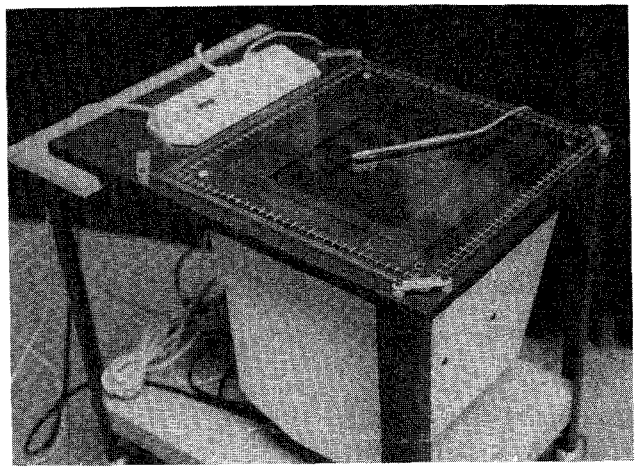

Figure 1: A Commune station.

observations of use, analysis, and prototypes. Our three-way Commune study was formulated based on our previous experiences with Commune, existing shared drawing tools, and known work on groups.

COMMUNE BACKGROUND

Commune presents each user with a horizontally-oriented monitor overlaid with a transparent digitizing tablet as shown in Figure 1. Users in distributed locations use styli to contribute to sketching activity on the screen by simply pointing or drawing directly on the tablet/monitor assembly. All of the participants may be active (marking, pointing, or erasing) simultaneously.

Commune is modelled after a pad of paper with markers, a medium we had seen a number of designers use in our previous observational work on face-to-face collaboration. This notion of the interface as a marker pad begins with the input technique–the user simply draws on the surface with a stylus–and extends to navigation–the user is presented with a sequence of sheets that may be revisited by paging back and forth. As the user moves the stylus, a pencil-shaped cursor tracks the motion; as a user presses down with the stylus, marks appear. Each user's marker (and "ink") is a distinct color so that even when users draw in the same area, their marks may be identified. When a user enters ERASE mode, the cursor is altered to reflect this change, and pressing down on the stylus causes marks to be erased rather than drawn.

The design of Commune has been based on studies of the use of traditional drawing surfaces. As mentioned previously, the initial studies have been followed by a development effort characterized by rapid prototyping, trial uses in representative situations, and additional studies of the use of the prototypes [Minneman and Bly, 1990]. We have found the following shared drawing system design requirements:

- Marks and gestures should be made visible to all participants without significant delay;
- Rapid switching among drawing, writing, and gesturing should be possible;
- Users should be able to mark, erase, and gesture in the same space simultaneously;
- Familiar mechanisms for drawing space activity should be maintained [Bly and Minneman, 1990].

In our studies of pairs of collaborators, we have been satisfied that Commune meets these goals and that our users have been able to do actual design work with the prototype. In addition, we have observed ways our users take advantage of capabilities not present in face-to-face drawing situations: a common orientation to the drawing surface and the ability to mark and point in exactly the same place at the same time. Our next step was to expand the use of Commune and observe whether it continued to support these interactions.

Related Shared Drawing Tools

Timbuktu, Remote Control, Carbon Copy, and related commercial products [Coleman, 1990] have been exploring the extent to which unmodified, single-user programs can be utilized in multi-party situations. While these programs have not investigated shared drawing per se, a drawing package running in one of these environments becomes a rudimentary shared drawing tool. However, these systems are plagued by their single-user roots; input from the participants is either explicitly handed around in a baton-passing manner or simply summed. Our previous studies have shown that these techniques are unnatural and inadequate for supporting interaction in shared drawing situations.

On the research front, a number of shared window system and dedicated multi-user application efforts have been directed at the programming and use issues involved in providing truly multi-user applications. The various tools, surveyed by Lauwers and Lantz [1990], each support some sort of shared drawing activity, ranging from simple bitmap sketching to moderately complex structured graphics. Another family of research projects have been utilizing video to provide users with novel ways of sharing a workspace. Several of Krueger's VideoPlace [1982] demonstrations have given pairs of users ways of drawing together by processing video images, extracting information about the locations of body features, and using those data as input to rudimentary drawing programs. The VideoDraw work [Tang and Minneman, 1990] has done away with the

computer entirely and provides its users with a shared sketching surface that uses video to convey the participants' hands and gestures as well as the sketched marks themselves.

In the majority of these cases the study of shared drawing activity itself has not been a major focus of the research. Conducting use studies of the sort presented here with these related tools would aid in evaluating the ramifications of making different design and implementation decisions than those embodied in Commune.

Studies of Distributed Group Work

Our intuitions in this study, shared by many who study small group behavior, were that "dyads are very different from larger groups" [Levine and Moreland, 1990]. However, while the literature base on group work is diverse and extensive, little work is available that directly addresses the effects of moving from dyadic to triadic groups in focused task performance. Foushee [1984] reports on changes in the social relations of two and three person aircraft flight crews, but we have not found any studies that report on this effect in distributed work situations.

There is an additional body of work on the effects of providing people with various communications media in distributed problem solving activity [McGrath, 1984; Chapanis, 1975]. We view our work here as an *extension* to that body of work in two distinct directions: first by extending the communication modes with the addition of a shared drawing tool and second by having our participants engage in a realistic design activity.

The Office Design Project [Weber and Minneman, 1987; Stults, 1989] looked at the work activity of three architects collaborating on a design task in a distributed setting supported by real-time audio and video connections. The video connections transmitted overhead shots of the drawing spaces as well as views of the designers' faces. The designers exhibited a high degree of engagement during the two-day session and made good progress on a difficult design program. Nevertheless, we observed several instances in which the designers were frustrated by not having a shared drawing space and attempted to compensate for that lack.

THE THREE-WAY COMMUNE STUDY

From the outset, our central concern in the Commune work has been to come to a better understanding of how to support collaborators *getting their work done*. Thus both Commune itself and the study setting were planned with the intent of providing a working environment where familiar ways of interacting would yield success and the task would engage the participants in naturalistic work activity. Our methodology is based on collecting real-time observations and videotape recordings of the participants' work activity and subjecting it to examination loosely based on interaction analysis [Suchman, 1987]. The current Commune study was not intended to verify or disprove any particular hypotheses; we do not believe that enough is known about this field at this time to adequately formulate such hypotheses or to design experiments to isolate the effects of a given hypothesis.

We entered the three-way Commune study expecting to see instances of concurrent activity, but unsure of what form those instances would assume or what purposes they might serve. With only two users, there is no confusion about who is contributing what pieces to the drawing surface (since, if it's not you, it's your collaborator). The negotiations for control of the discussion are simpler and are thus less affected by the distributed nature of the environment. With three users, we expected to find more confusion in proceeding with the task, both in sketching activity and in discussions of the problem. We were not sure that 3-way Commune would support behaviors we'd seen with two users: the ways that participants time their talking, marking and gesturing activities; participants' uses of the ability to be in the same place at the same time; the effect of cursors as a substitute for gesture; the relationship of the shared drawing surface and the video/audio links; and methods that the participants use to demonstrate that they are following the current discussion.

For the exercise described in this paper, the Commune hardware was configured as shown in Figure 2. Commune workstations were set up in three office spaces equipped with audio and video connections; a typical setup for the exercise is shown in Figure 3. Note the horizontal surface of the Commune shared drawing system, the two monitors displaying images of the other two participants, and the head-on camera for capturing the view being sent to the other two collaborators.

Subjects for the Commune exercise were solicited from the PARC community via a broadly-distributed email message. The message offered two alternatives for participating in the exercise: come with an established group and bring your own task or be randomly assigned to a group to work on a problem of our creation. We scheduled eight sessions: five that would participate in our design exercise and three that would use the Commune setting for one to two hour sessions with their own work. This paper reports on our observations of the five groups participating in the assigned task.

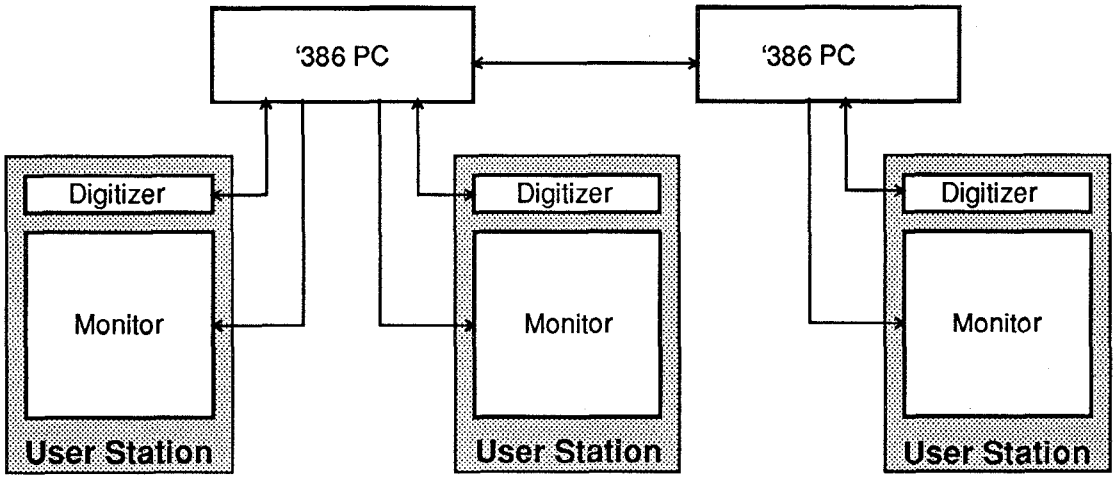

Figure 2: Block diagram of the system as configured for this study: three user stations, two PCs.

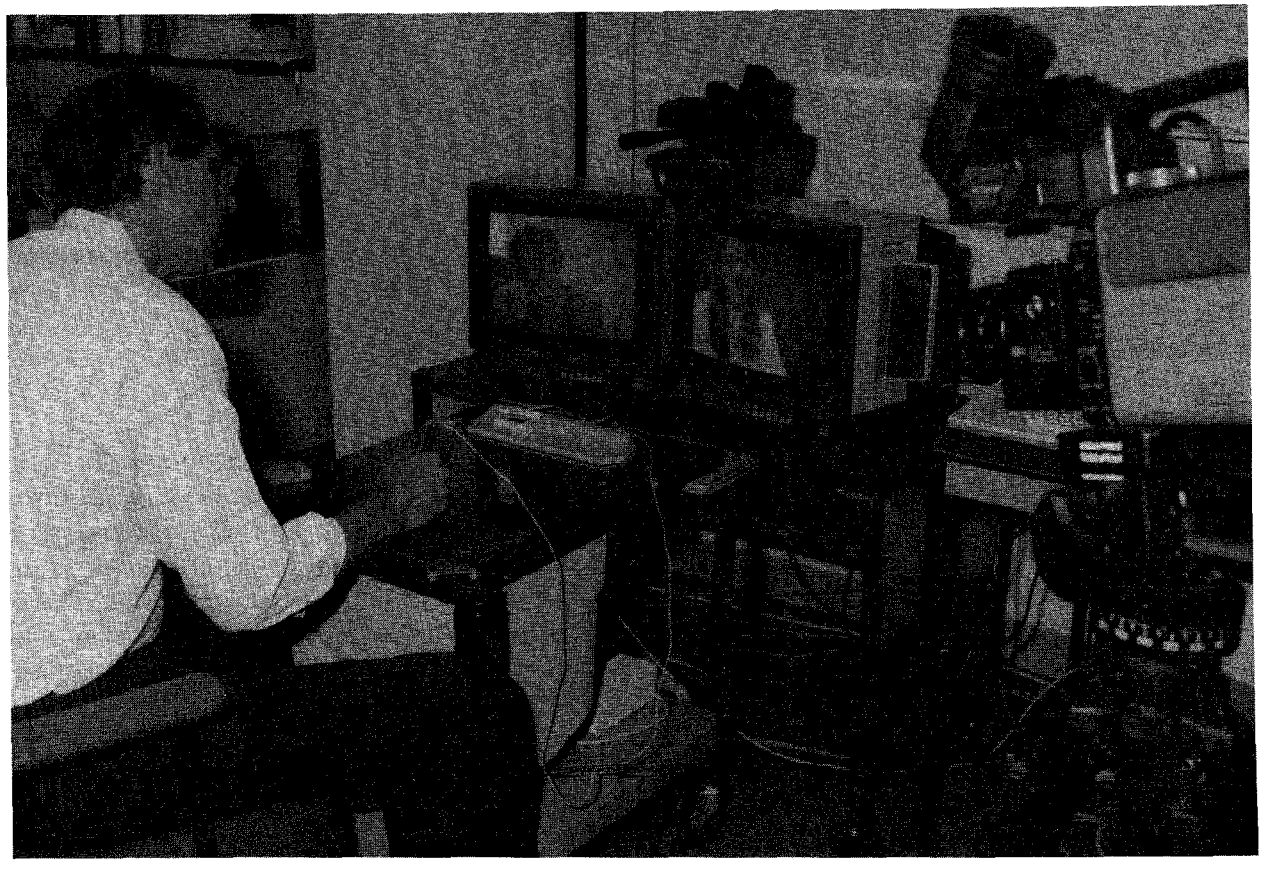

Figure 3: A typical setup for the study: a Commune station and audio/video link equipment.

Our participants represented a variety of job and experience backgrounds, and we arbitrarily formed threesomes. The assigned exercise was a brief design task involving the user interface of a credit card gasoline pump. A copy of the actual handout given to the participants is included as an appendix to this paper. While this task was obviously not "real work" for the participants, they rapidly became engaged in the task and it quickly became their primary focus (as opposed to the peculiarities of their working environment). In addition, while the participants clearly had little stake in the outcome

of their work, they did exhibit concern about the quality of their efforts and often discussed the quality of their solutions after their session was over. Of the five groups that were subjects for this exercise, three had Commune connections augmented by video and audio connections; the remaining two groups had Commune and audio connections only.

Subjects would arrive in a common area (and, if necessary, meet one another), be taken to their respective working rooms, and given a brief (~90 second) introduction to Commune and the rest of their collaboration environment. They were then presented with the problem statement and left to work on their own for 20-25 minutes. At the end of that period they were asked to take a couple minutes to summarize their design activity and then to spend a few additional minutes discussing their impressions of working in the design environment we had provided.

A comprehensive data set was collected for each working session. Included were videotapes, hardcopy images of the participants' Commune screens, complete logs of all of the Commune activity (which may be replayed if needed for future analysis), and notes we made in real time about significant events in the work. The five videotape recordings made during each session consist of an over-the-shoulder view of each of the three participants' workspaces, a four-into-one videotape recording of an unobstructed Commune screen plus the three face views being transmitted among the participants, and a full-screen view of a the Commune screen.

RESULTS

At this time, we have performed only a preliminary examination the data from the gas pump interface design sessions to form an initial set of observations and to raise issues that might be considered in a more detailed and focused analysis. Our previous observations of drawing activity led us to believe that drawing space activities are as important as the resulting marks. In the study reported here, we again found that participants regularly combine marking activities with talk, that they rapidly move among drawing, writing, and gesturing, that they interact on the same drawing marks, and that they use marks and gestures to illustrate and reference ideas.

In all cases, our subjects were able to complete an initial design for the gasoline pump that met the written requirements. Figure 4 is a sample page from one of the sessions. The problem statement presented the participants with a familiar situation, and they quickly became engaged in design activity. None of the groups reported or demonstrated any difficulty in understanding the problem or figuring out how to proceed.

Three Users vs.Two Users

Our primary goal in this particular exercise was to uncover any difficulties that would arise with three Commune users rather than two. Our hope was that we could observe ways in which the shared drawing technology either contributed to any difficulties or could be modified to support easier interactions.

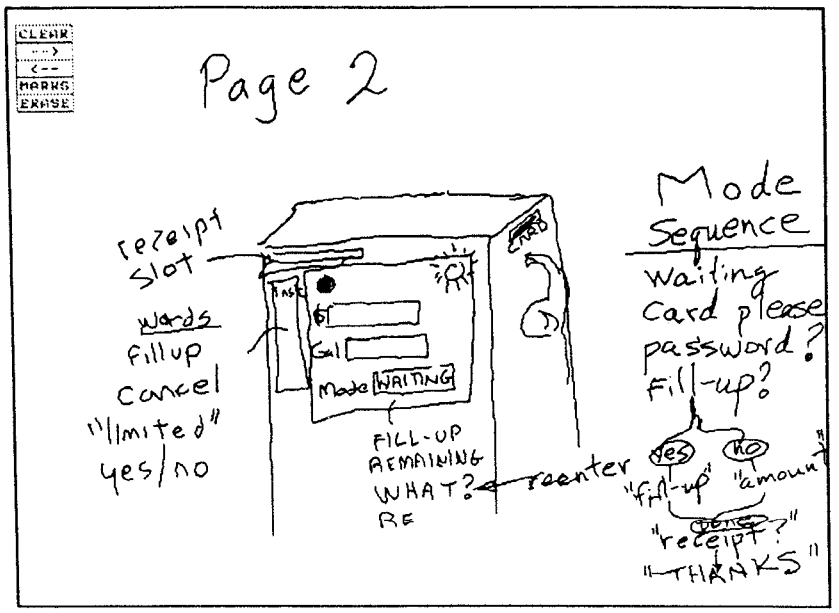

Figure 4: A sample page of shared drawing marks.

In fact, we found no significant breakdowns in the distributed design tasks with three participants. At the beginning of the sessions, there was generally talk about knowing which cursor was in use by which of the subjects. However, after a group started working on the task, the accompanying talk seemed to resolve any ambiguity. There were very few problems involving confusion about the identity of the participant controlling a particular cursor or making a particular mark. As expected, the situation of one group member playing a particularly strong role did in fact occasionally arise. However, again the Commune and audio or audio/video connections supported typical participant interactions.

Video/Audio Connections vs. Audio-only

In many respects, the two groups that had audio-only connections (in addition to Commune) performed much like the three groups that had both video and audio connections. They became engaged in the task, used Commune with equal facility, and created designs that were similar to those in the other groups. However, our initial observations of the videotapes have suggested the following about the audio-only groups:

- Although individuals continue to look up from the drawing surface as though they're seeing someone, they look up less frequently and the motions are not in sync with other participants' gazes.
- Participants tend to engage in meta-level discussions less often.
- Participants use gestures in the air less often.
- Participants made more assumptions about whether there was agreement on issues and whether all members were ready to proceed.
- Participants became visibly detached from the problem activity occasionally (e.g., staring off and then noticeably reorienting).

Passive Participation

Perhaps the most interesting observation from these studies, thus far, comes from the way in which participants used Commune while someone else was speaking. There were not only times when participants worked on different segments of the task but other non-verbal activities occurred as well. These included erasing past work that was no longer under consideration (e.g., a zoomed-in view used to explain a particular detail would be erased after it has served its purpose), writing down ideas that had not been taken up by the group (e.g., that the new gas pumps should also dispense cash), and adding details to existing sketches (e.g., drawing the hoses on a gas pump sketch). We also saw many instances where all three participants would simultaneously contribute to the same drawing activity (e.g., collaboratively drawing a sun shade over the amount due display while discussing its utility).

DISCUSSION

We believe our most significant observation is the lack of any major surprises in the transition from two to three participants. Commune continues to meet our expectations in supporting distributed drawing activities, and although there were certainly differences in the social relations of the groups, there were no breakdowns that could be attributed to the shared drawing tool. Our users continued to take advantage of the ability to intermix talking, marking, and gesturing rapidly and to interact together in the same spaces at the same time; our shared drawing tool design requirements stand unchanged.

One topic area where we expected that our participants might experience problems with the current Commune implementation was the identification of users with their cursors and marks. One of our intuitions about design changes that might be suggested by this study was that of additional cues about mark and cursor identity; this was not the case. Recall that each of the three Commune participants had a different color cursor and left different color marks. Some participants did appear to learn the correspondence among people and colors. (For a couple of the sessions, we put colored stickers on the monitors with the face view corresponding to the person's cursor color. This explicit mapping between a person and their cursor color did not appear to be referenced.) We speculate that there was not much confusion because it is easy to identify what is being done in concert with talk (and if it's not in concert with talk then identity isn't as important at that moment).

On first glance, our study appears to support earlier work that suggests that the presence of video does not significantly affect group task performance. However, we believe our data contains evidence that the video connection does make a difference. For example, colleagues watching our analysis tapes of the face views and Commune screen had no trouble identifying which sessions included video connections and which did not. We expect that these differences are subtle and most likely to become meaningful in long-term distributed use situations. Many of the clues we observe do not appear to effect the short term tasks but will likely lead to disengagement from the task over longer time periods. The gazes off into the distance, the lack of response to many suggestions and ideas, the tendency not to engage in meta-

discussions all suggest that the participants are not as involved in the task as they are when video connections are available.

We are particularly interested in the many instances of non-verbal or passive participation. A drawing tool such as Commune that provides a shared surface with a common orientation to all participants seems to be particularly suited for supporting non-verbal interactions in the task activity. Private space that is common is most face-to-face shared drawing situations becomes part of the immediately visible public space in Commune. Thus, the activities and contributions of all participants are an integral part of the on-going interactions. Rather than shy away from interaction by doodling on a separate sheet of paper or becoming disengaged in the task, participants seemed to make use of the ability to become involved in the task activity (e.g., after having an idea rejected in an unsatisfactory manner, a participant can jot that idea onto the shared drawing surface where it will be rediscovered for later consideration).

CONCLUSIONS

With three users, Commune continues to support the drawing activity of distributed design groups in ways similar to the support of two persons. However, our recent exercise raises some interesting issues that we believe deserve further consideration:

- Can a shared drawing surface offer effective new opportunities for individual participation in collaborative activities?

- Although tasks continue to be completed by groups having no video connection to one another, will long-term studies support our observations that participants become less-engaged in the task and that over time, task output would not be as productive?

Our next steps for Commune will be to install the prototype in a longer-term real use situation. We feel comfortable that it offers good support for designers who must work at a distance from one another. At the same time, we are particularly interested in more detailed analysis of the ways in which participants appear to use Commune capabilities that are not offered by traditional shared drawing surfaces. Studies to observe the effects of long-term video/audio versus audio-only communication links will rely on more concentrated attention to the group interactions as participants engage in all their working activity, not just shared drawing.

ACKNOWLEDGEMENTS

We thank all of our volunteers for their willingness to try something new so that we could learn something new. We also thank our colleagues in the Design, Use, and Shared Spaces group who gave not only moral support but space and equipment as well.

REFERENCES

Bly, Sara A., and Scott L. Minneman, "Commune: A Shared Drawing Surface", *Proc. of the Conference on Office Information Systems* (Boston, MA, April 1990), pp. 184-192.

Chapanis, Alphonse, "Interactive Human Communication", *Scientific American*, Vol. 232, No. 3, (March 1975), pp. 36-42.

Coleman, Dale, "Timbuktu vs. Carbon Copy Mac: Close Race", *MacWEEK*, (September 11, 1990), pp. 66-68.

Foushee, H. Clayton, "Dyads and Triads at 35,000 Feet: Factors Affecting Group Process and Aircrew Performance", *American Psychologist*, Vol. 39, No. 8 (August 1984), pp. 885-893.

Krueger, Myron W., *Artificial Reality*, Reading, MA: Addison-Wesley, 1982.

Lauwers, J. Chris, and Keith A. Lantz, "Collaboration Awareness in Support of Collaboration Transparency: Requirements for the Next Generation of Shared Window Systems", *Proc. of the Conference on Computer Human Interaction (CHI) '90* (Seattle, WA, April 1990), pp. 303-312.

Levine, John M., and Richard L. Moreland, "Progress in Small Group Research", *Annual Review of Psychology*, Vol. 41 (1990), pp. 585-634.

McGrath, Joseph E., *Groups: Interaction and Performance*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1984.

Minneman, Scott L., and Sara A. Bly, "Experiences in the Development of a Multi-user Drawing Tool", *Proc. of the Third Guelph Symposium on Computer Mediated Communication* (Guelph, Ontario, May 1990), pp. 154-167.

Stults, Robert, *Experimental Uses of Video to Support Design Activities*, Xerox PARC Technical Report #SSL-89-19, December 1988.

Suchman, Lucy A., *Plans and Situated Actions: The problems of human-machine communication*, Cambridge, UK: Cambridge university Press, 1987.

Tang, John C., and Scott L. Minneman, "VideoDraw: A Video Interface for Collaborative Drawing", *Proc. of the Conference on Computer Human Interaction (CHI) '90* (Seattle, WA, April 1990), pp. 313-320.

Tatar, Deborah G., "Using Video-Based Observation to Shape the Design of a New Technology", *SIGCHI Bulletin*, Vol. 21, No. 2 (October 1989), pp. 108-111.

Weber, Karon, and Scott L. Minneman, "The Office Design Project", (videotape), Palo Alto, CA: Xerox PARC, 1987.

APPENDIX

A Design Problem:

> Green Oil has long been dedicated to providing its customers with the utmost in quick, convenient service. A recent satisfaction survey has indicated that patrons are unhappy with the amount of time required to pay the cashier in self-serve stations, especially in urban areas where filling the tank requires two separate trips to the clerk.
>
> A suggested solution to this source of customer dissatisfaction has been to install credit card gas pumps in all of the chain's stations. Using a credit card pump, a user would charge the gas using a credit or bank card. Green Oil executives are very excited about this concept and hope to have prototypes in the field within a year.
>
> Your task is to provide Green Oil with a design for a new credit card pump. The set-up might look much like an automatic bank teller machine, though you should feel free to do as you wish about this issue. Keep the environment in mind; not all technologies are robust enough to be used in gas station pumps.
>
> Please create a design for the user interface; be sure to include options both for fill-ups and for specific dollar amounts. Plan to spend no more than 30 minutes on the task. Thanks!

When you feel satisfied with your ideas for Green Oil's credit card pumps, would you all spend a few minutes talking about the experience of working together with the shared drawing stations? Thank you.