

The Effects of Workspace Awareness Support on the Usability of Real-Time Distributed Groupware

CARL GUTWIN

University of Saskatchewan

and

SAUL GREENBERG

University of Calgary

Real-time collaboration in current distributed groupware workspaces is often an awkward and clumsy process. We hypothesize that better support for *workspace awareness*—the understanding of who is in the workspace, where they are working, and what they are doing—can improve the usability of these shared computational workspaces. We conducted an experiment that compared people's performance on two versions of a groupware interface. The interfaces used workspace miniatures to provide different levels of support for workspace awareness. The basic miniature showed information only about the local user, and the enhanced miniature showed the location and activity of other people in the workspace as well. We examined five aspects of groupware usability: task completion times, communication efficiency, the participants' perceived-effort, overall preference, and strategy use. In two of three task types tested, completion times were lower in the awareness-enhanced system, and in one task type, communication was more efficient. The additional awareness information also allowed people to use different and more effective strategies to complete the tasks. Participants greatly preferred the awareness-enhanced system. The study provides empirical evidence that support for workspace awareness improves the usability of groupware, and uncovers some of the reasons underlying this improvement.

Categories and Subject Descriptors: D.2.2 [**Software Engineering**]: Design Tools and Techniques—*User interfaces*; D.2.8 [**Software Engineering**]: Metrics—*Performance measures*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology*; H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Synchronous interaction*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*Interaction techniques*

This research was supported in part by the Natural Sciences and Engineering Research Council of Canada, and in part by Intel Corporation. This article is a substantially expanded version of a report presented at the 1998 ACM CHI conference [Gutwin and Greenberg 1998]. Authors' addresses: C. Gutwin, Department of Computer Science, University of Saskatchewan, 57 Campus Drive, Saskatoon, Saskatchewan S7N 5A9, Canada; email: gutwin@cs.usask.ca; S. Greenberg, Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary, Alberta T2N 1N4, Canada; email: saul@cpsc.ucalgary.ca. Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 1073-0516/99/0900-0243 \$5.00

General Terms: Experimentation, Human Factors, Measurement

Additional Key Words and Phrases: Computer-supported cooperative work, real-time distributed groupware, usability, workspace awareness

1. INTRODUCTION

Real-time distributed groupware allows people to work together at the same time from different places (e.g., Baecker [1993] and Greenberg [1991]). Many of these systems provide shared computational workspaces—two-dimensional areas akin to whiteboards or tabletops—where people can create and manipulate task artifacts. Although many of the technical problems of constructing these systems have been solved, their usability problems have not yet been eliminated. Collaboration in groupware workspaces is often awkward, stilted, and frustrating compared to collaboration in face-to-face settings. The difficulty is particularly acute when the workspace is larger than the screen and people navigate independently through the workspace—called relaxed-WYSIWIS view sharing [Stefik et al. 1987].

Part of the problem with current systems is that they do not provide much information about other participants in the session. When people work together in a face-to-face setting, a wide variety of perceptual cues help them keep track of what others are doing. This awareness of others in the workspace is *workspace awareness*, the up-to-the-moment understanding of another person's interaction with the shared space [Gutwin 1997; Gutwin and Greenberg 1996]. At a simple level, it involves knowledge of who is present, where they are working, and what they are doing. Workspace awareness is used in collaboration to coordinate activity, to simplify verbal communication, to provide appropriate assistance, and to manage movement between individual and shared work.

We believe that being able to maintain workspace awareness is necessary for natural and smooth collaboration in a shared workspace. Current groupware systems, however, provide only a fraction of the information needed to maintain workspace awareness. Consequently, we hypothesize that increased support for workspace awareness will improve the usability of real-time distributed groupware. Our goal in this article is to evaluate that hypothesis, and in what follows we describe a study that we carried out to assess the effects of workspace awareness support on a realistic groupware system.

A previous investigation provided qualitative evidence that awareness support is valuable [Gutwin et al. 1996]. In particular, it showed that *workspace miniatures*—miniature representations of the entire workspace—are useful tools for any kind of work done in medium-sized workspaces, and are useful vehicles for awareness information. That study convinced us of the value of workspace miniatures in general; however, we wished to focus more closely on the usefulness of providing awareness information in the miniature. Therefore, the current study focuses on the

quantitative effects of awareness support on groupware usability. We compared two groupware interfaces that provide different amounts of awareness information through their workspace miniatures. In particular, we compared a basic miniature to one that adds three kinds of information:

- the location of others' viewports in the workspace;
- the location and motion of people's cursors; and
- the motion of workspace objects as they are moved.

The awareness-enhanced version of the miniature is called the *radar view*, a groupware device first described as part of the SharedARK system [Smith 1992; Smith et al. 1989]. Our experiment measured three general aspects of groupware usability—how well groups perform with each interface, the efficiency of their collaboration, and the group's satisfaction with the system—and looked at the strategies that the two different groups used. Between-participants tests showed that the extra awareness information allowed for significant improvements in speed and communication efficiency for some of the tasks, and within-participants tests showed improvement on all tasks in speed and perception of effort. Analysis of strategy suggests that the radar view allowed people to use more effective strategies for completing tasks. In addition, when the awareness information did not improve performance, there were no cases where it significantly impaired performance. Observations of the sessions and participant feedback provide some explanation for these results, and indicate additional design directions. Before describing the experiment and our findings, we begin by outlining the two basic ideas underlying the research: workspace awareness and groupware usability.

2. WORKSPACE AWARENESS

It is becoming more and more apparent that being able to stay aware of others plays an important role in the fluidity and naturalness of collaboration (e.g., Dourish and Bellotti [1992], Segal [1995], Tang [1991], and Gaver [1991]). We have looked closely at one kind of awareness that involves collaboration in a shared workspace. *Workspace awareness* is the up-to-the-moment understanding of another person's interaction with a shared workspace [Gutwin 1997; Gutwin and Greenberg 1996]. It involves several kinds of knowledge about the people in the workspace and their activities. This knowledge is useful for many of the activities of collaboration—for coordinating action, managing coupling, talking about the task, anticipating others' actions, and finding opportunities to assist one another. Since workspace awareness information is dynamic, maintaining it entails a continuous process of gathering information from the workspace environment and integrating that information with existing knowledge.

We have built a conceptual framework of workspace awareness that sets out its component elements, mechanisms for maintaining it, and the ways that workspace awareness is used [Gutwin 1997; Gutwin and Greenberg

Table I. Elements of Workspace Awareness Relating to Real-Time Activity

Category	Element	Specific Questions
Who	Presence	Is anyone in the workspace?
	Identity	Who is participating? Who is that?
	Authorship	Who is doing that?
What	Action	What are they doing?
	Intention	What goal is that action part of?
	Artifact	What object are they working on?
Where	Location	Where are they working?
	Gaze	Where are they looking?
	View	Where can they see?
	Reach	Where can they reach?
When	Event history	When did that event happen?
How	Action History	How did that operation happen?
	Artifact history	How did this artifact come to be in this state?

1999]. The framework is built from a synthesis of existing work on awareness (both in CSCW and in Human Factors) as well as from our own observational studies of group work (see Gutwin [1997]). The first part of the framework divides the concept of workspace awareness into components. People keep track of many things in a shared workspace, but certain elements make repeated appearances in research literature (e.g., Dourish and Bellotti [1992], Sohlenkamp and Chwelos [1994], and McDaniel and Brinck [1997]). This basic set answers “who, what, where, when, and how” questions. That is, when we work with others in a shared space, we know who we are working with, what they are doing, where they are working, when various events happen, and how those events occur. People keep track of these things in all kinds of collaborative work, and these are the kinds of information that should be considered first by designers. In the current study, however, we are particularly interested in the elements of workspace awareness that relate to real-time activity; in Table I, these are the first three categories of Who, What, and Where.

When people know the answers to these questions, many of the activities of collaboration are made easier. For example, knowing where another person is working and what part of the workspace they can see allows for efficient means of communication such as deictic reference (e.g., Tatar et al. [1991]). Pointing to an object is much easier than indicating it through description, but deixis requires that people know what is visible to the other person. Coordination is another example. Coordinating actions with another person in a shared workspace is far simpler when both parties know what the other is doing. Workspace awareness is particularly evident in continuous activities where people work with the same objects (e.g., Tang [1991]).

The awareness problem in groupware is that while it is relatively easy to answer the questions of Table I in a face-to-face workspace, maintaining awareness is much more difficult in a distributed workspace. It is often difficult or impossible to keep track of others in a groupware system, because groupware provides only a fraction of the perceptual information

that is available in a face-to-face workspace. The overall approach in our research is to recreate some of the awareness information that is missing from a groupware workspace, allowing people to gather and use it just as they do in the real world. As might be expected, there are several issues to be resolved in adding awareness information to a groupware system, such as what information to add, how to present it in the interface, and when to make it available. Although we will not discuss those issues here, we will note that the awareness displays used in this experiment have undergone several design revisions to address these questions (see Gutwin [1997]).

3. GROUPWARE USABILITY

Our hypothesis is that awareness support will improve groupware usability. Testing this claim implies knowing what groupware usability is, and knowing how to measure it. Since no concrete definition of groupware usability has been accepted by the CSCW community, we adapt the concept from the better-known area of “singleware” usability. Usability in a single-user environment is the degree to which a system is effective, efficient, and pleasant to use, given a certain set of users and tasks (e.g., Nielsen [1993] and Shackel [1990]). Real-time groupware systems are subject to these criteria as well, but now two kinds of activity must be considered: taskwork and teamwork. Taskwork is the domain activity, the activity that produces things like drawings, documents, or models. A groupware system clearly must allow taskwork to proceed effectively, efficiently, and pleasantly, in order to be a good application. However, groupware must also support teamwork—the work of working together—in order to be truly usable.

Teamwork involves several activities: for example, group members must communicate, organize joint action, provide assistance, coordinate activity, divide labor, and monitor each other’s work. Each of these activities can be considered in terms of efficiency, effectiveness, and participant satisfaction. Our conception of groupware usability includes both taskwork and teamwork, and we define it as the degree to which a groupware system supports the activities of collaboration. Taskwork is represented in the activity of execution, and teamwork is represented by those activities that accomplish the mechanics of collaboration: communication, coordination, planning, monitoring, and assistance. Although teamwork also includes social and affective elements, we limit our definition of groupware usability because we believe that the features or facilities in a groupware system have a much greater likelihood of affecting the group’s communication, coordination, or execution than they do of affecting social issues and group dynamics.

The second issue in evaluating our hypothesis is one of measurement: how can improvements in usability be determined? Groupware is notoriously difficult to measure (e.g., Grudin [1990]); the main problem is that usability, effectiveness, efficiency, and pleasantness are qualities that cannot be directly observed. Other researchers, however, have found indirect measures that appear to fit well with our conception of groupware

usability (e.g., Olson et al. [1992; 1995]). In particular, these studies measure three aspects of collaboration: product, process, and satisfaction.

- Product measures* relate to the taskwork elements in groupware usability, and measure aspects of the outcome of a task. Product measures are based on the assumption that the usability of a system will influence the group's success in completing the domain task. They judge collaboration outcomes in terms of quality or time.
- Process measures* relate to the teamwork elements of groupware usability. They look for patterns in behavioral or verbal activity during a collaborative session and connect these to issues of effectiveness or efficiency. Process measures are often obtained by observation or through video and audio protocol analysis.
- Satisfaction measures* can relate to both teamwork and taskwork, but focus on the participant's own beliefs about the activities of collaboration. These measures record participants' subjective experience with the groupware system, and whether the participants find the system to be a good setting for collaboration. Satisfaction measures are usually obtained with questionnaires and interviews.

Product measures have been widely used in those evaluations that follow the psychological tradition, and assess either quality of product using expert judgment or task completion time (e.g., Olson et al. [1992; 1995]). Several studies have looked at how different communication facilities affect the time it takes to complete a group task. For example, Chapanis [1975] showed that a distributed construction task was accomplished more quickly when participants could talk to one another than when they communicated by either handwriting or typing, but that the addition of a video link did not have any effect. One of the main problems with product measures in problem-solving tasks, however, is that people are able to “work around” many kinds of difficulties; as a result, product measures are “only sensitive to gross changes in the facilities available for communication” [Monk et al. 1996, p. 125]. Consequently, many studies of product measures report no differences (see summary in Anderson et al. [1997]).

Process measures have primarily been used to explore communicative and task activity. Unlike product measures, these have proved much more successful for observing differences between different media channels and for understanding the nuances of how people interact over these channels. In many cases, techniques from Conversation Analysis are used to quantify and compare various aspects of conversational behavior, such as the amount of time spent on particular topics (e.g., Olson et al. [1995]), the linguistic types of the utterances (e.g., Tatar et al. [1991] and Monk et al. [1996]), or communicative breakdowns and repairs (e.g., Doerry [1997]). For example, Olson and colleagues [Olson et al. 1992] look at the transcripts of collaborative sessions to determine how extensively different groups explored the design space for a problem-solving task (the design of an automated post office). A variety of researchers have quantified the

effects of video-mediated communication on process. O’Conaill and Whittaker [1997] found that measurable characteristics of speech processes such as backchannel communication, interruptions, overlaps, and turn-taking were effective in examining face-to-face interactions. Anderson et al. [1997] looked at dialog length in terms of turns and words, how turn-taking is managed, and nonverbal communication such as how gaze is used. Heath et al. [1997] examine how individuals coordinate their talk and activities: for example, one person looking at another typically signals that the person is available and prepared to listen; speech devices (such as pauses and sound elongation) encourage one person to turn toward the other; and hand gestures invite more active participation by the other person. Other examples of process measures assess where people look during a collaborative interaction [Gaver et al. 1991], when they make eye contact [Ishii et al. 1992], or when they use communicative gestures [Tang 1991].

Satisfaction and preference measures are primarily used in combination with other measures or with observational techniques to determine how participants feel about product and process issues (e.g., Gutwin et al. [1996] and Olson et al. [1992]). Satisfaction measures are often more sensitive to certain factors than objective assessments (e.g., Olson et al. [1992] and Whittaker [1996]): for example, several studies have shown that people like having video of their collaborators (e.g., Anderson et al. [1997] and Olson et al. [1997]), even though product and process measures are unable to ascertain any difference between conditions that have video and those that have only audio (e.g., Egido [1990]).

Evaluations that consider product, process, and satisfaction aspects of groupware usability can provide a broad and balanced look at a groupware system. This is the approach that we follow in our experiment. Below, we review our methods, and discuss each of the measures we use in these three areas.

4. EXPERIMENTAL METHODS

The experiment tests the hypothesis that increased workspace awareness support increases groupware usability. We compared people’s collaboration when using two groupware interfaces—each providing different amounts of awareness information through workspace miniatures. In this section we outline the groupware system used, the experimental tasks, the participants, the study design, and the measures taken.

4.1 System and Experimental Conditions

We are interested in groupware systems that allow small groups to collaborate in real time in a medium-sized visual workspace. Activities in these systems are organized around the creation, manipulation, and organization of artifacts in the workspace. We built such a system for this experiment, using the GroupKit toolkit [Roseman and Greenberg 1996] and the Pad++ drawing system [Bederson and Hollan 1994]. The application was a pipeline construction kit that allows the assembly and manipulation of simple




Fig. 1. The pipeline application (radar view version).

pipeline networks in a shared two-dimensional workspace (Figure 1). Users can create, move, and rotate sections of pipe, and can join or split sections using a welding tool. The workspace is rectangular, and four times larger than the computer screen in each direction. Users scroll around the workspace by dragging their cursor past the window border.

The pipeline system's interface consists of two windows. The *main view* takes up most of the screen and shows objects in full size and detail. The main view allows users to manipulate objects and to scroll to other areas of the workspace. People create pipelines by dragging pipe sections from storehouses in the corners of the workspace (see Figure 1), aligning the sections, and then welding them together by dropping a diamond-shaped welding tool onto the joint. Welds are marked by a yellow square, and once pieces are welded, they move as a unit.

The second window is one of two miniature views, the *radar view* or the *overview*. This view is inset into the top left corner of the main view, and shows the entire workspace in miniature. The radar view and the overview differed in three ways, as compared in Figure 2.

- (1) *Update granularity*: The radar showed workspace objects as they moved; the overview was only updated after the move was complete.




Fig. 2. Radar view (left) and overview (right).

- (2) *Viewport visibility*: The radar showed both people's viewports (the area of the workspace visible in each person's main view), and the overview showed only the local user's viewport.
- (3) *Telepointer visibility*: The radar showed miniature telepointers for both users, and the overview did not show any telepointers.

In sum, the two conditions differed only in the awareness information presented in the miniature. The overview only showed information about the local user, while the radar showed where the other person was located, showed their pointer, and showed moves as they occurred. In terms of the categories of workspace awareness shown in Table I, the radar view provides a variety of information that can be used to answer Who, What, and Where questions.

The way we have set up the experimental conditions deliberately explores three variables at once, but we believe that the three kinds of awareness information listed above should be studied as a complete package. In some cases, the different kinds of information must be provided together in order to make sense to the viewer (for example, telepointers require fine-grained object updates if object dragging is to be understandable). In addition, our previous qualitative studies with the radar view [Gutwin et al. 1996] have indicated that these kinds of information work well together. In the current study, although we group the three variables for quantitative analysis, we use our observations of the way people performed the tasks to provide some indication of which type of information is important for a particular task.

4.2 Tasks

Participants completed three different tasks. The test tasks were designed to mimic joint actions that we had previously seen in observations of




Fig. 3. Workspace map for the first Follow task.

face-to-face collaboration [Gutwin 1997], but were constrained to meet three criteria. First, we wanted tasks that required people to move independently around the workspace. Second, we wanted people to need awareness of their partner, and in particular, to use location or activity information. Third, we wanted tasks that were as realistic as possible and likely to occur in a wide variety of workspace activities. The test tasks are quite constrained; however, they are based on real and common joint actions. In the tasks below, division of responsibility is similar to Chapanis' [1975] communication studies, where a *source* person has information that a *seeker* person needs for his or her part of the task. These two roles will occur whenever people have different knowledge about the workspace and the task, and where one person must direct traffic, at least for a short while, and oversee the completion of some shared activity. The three tasks were called Follow, Copy, and Direct and are described briefly below (additional details on task instructions and materials are given in Gutwin [1997]).

4.2.1 Follow. The Follow task involved meeting another person at a specified location. Participants were asked to make 10 specific welds on an existing pipe network. One person, the joiner, was given a paper map (Figure 3) showing the locations to be welded, and had to prepare the pipe




Fig. 4. Initial workspace state (left) and first goal (right) for the first Copy task.

sections at each place. The other person was the welder, and would follow the joiner to each location and weld the pipe. Since the welder had no map, the joiner was also responsible for ensuring that the welder went to the correct location. The workspace map for the first Follow task, showing the initial pipeline layout and the 10 welding sites, is shown in Figure 3. The map also shows the initial state of the workspace at the start of the task.

The Follow task represents the situation of meeting another person at a specified (and not already well-known) location in the workspace. As a realistic example, consider two engineers discussing a schematic diagram. One engineer wishes to show the other how a particular component has been laid out, and must therefore arrange for both of them to navigate to the part of the workspace where the component is located. In the Follow task, we expect that awareness information about people's locations in the radar view will enable them to find each other more quickly and with less discussion.

4.2.2 Copy. The Copy task involves indicating objects to another person. Participants were asked to construct two identical structures from two existing stockpiles of pipe sections. The stockpiles were located at opposite ends of the workspace. One person, the leader, had a paper picture of what was to be built, and used this to find the next piece in the stockpile. The other person, the copier, did not have the picture, and so had to copy the leader's actions. The leader was responsible for making sure that the copier knew which piece to take next and where to place it. The initial state of the workspace and the first picture of what was to be built are shown in Figure 4.

The Copy task represents the situation of having to indicate an object to another person in the workspace. Indicating and deictic reference are ubiquitous in shared workspaces (e.g., Tang [1991] and Tatar et al. [1991]). We expect that faster update granularity and telepointers will allow people to see more quickly which object the leader has selected, and should therefore reduce the amount of verbal description necessary.

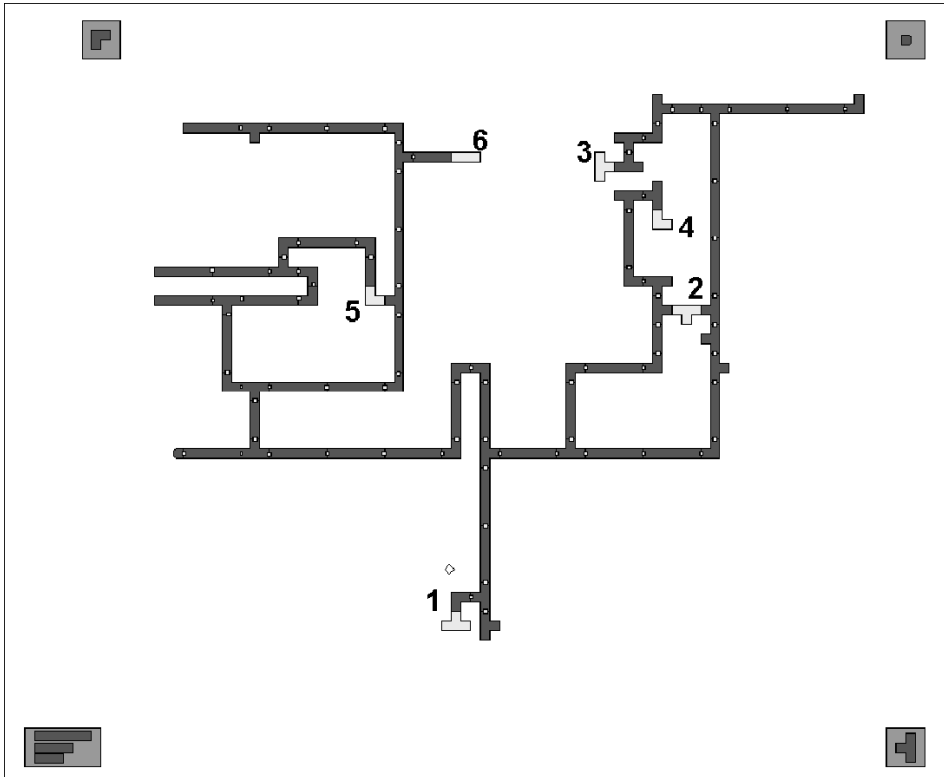


Fig. 5. Workspace map for the first Direct task.

4.2.3 Direct. The Direct task involves giving workspace directions. One participant was asked to verbally guide the other through adding six specific pipe sections to an existing network. The director had a map showing which pieces were to be added, and where they were to be added, but was not allowed to move around in the workspace. The actor did the work, following the director's instructions. The director did not see the main view during this task, so the only visual feedback that received of the actor's progress was from the miniature view. The workspace map for the first directing task is shown in Figure 5; the pieces to be added are shown in gray.

The Direct task represents the situation of telling another person how to get to a specific location; it is similar to the Follow task, but in this case the director does not go to the location (as done in Follow). This situation occurs whenever one person knows the location of a particular object, while another person does not (but needs to go there). In the Direct task, we expect that update granularity and telepointer movement will assist the director in understanding the current precise location of the actor, which will allow for more timely and more appropriate directions.

Table II. Experiment design (P = Pair). Arrows indicate comparisons.

		Task	
		Follow 1, Copy 1, Direct 1	Follow 2, Copy 2, Direct 2
View	Radar	P 1–10	↔
	Overview	P 11–20	

4.3 Study Design

The study used a complex design in which both between-participants and within-participants hypotheses were tested. The design and the comparisons made in the study are summarized in Table II. The design combines two independent variables in a two-way mixed factorial design: View is both a between-participants factor and a within-participants factor, and Task is a repeated-measures factor. The overall research hypothesis is that the additional awareness information in the radar view will improve the usability of a groupware system, where usability is measured in terms of people's speed, efficiency, and satisfaction. The specific hypotheses are as follows:

- (1) *Between-participants*: groups in the radar condition will complete the first three tasks more quickly, with greater efficiency and with greater satisfaction.
- (2) *Within-participants*: groups who use the overview first and then the radar view will have a greater improvement in speed and perception of effort than groups who use the radar view first and then the overview.

The hypotheses are tested by looking for effects of View (either main effects or effects in interaction with Task), using three dependent variables (completion time, verbal efficiency, and perceived effort). Effects of Task will not be investigated, since the tasks are considerably different and since large contrasts are expected.

Groups used both the overview and the radar view, so both interface order and task order were counterbalanced. Interface order was counterbalanced by having half the groups start with the overview, while the other half start with the radar view. Task order was counterbalanced by having groups carry out the three test tasks in different orders: group 1 did the tasks in the order Follow-Copy-Direct, group 2 in the order Follow-Direct-Copy, etc. The second set of tasks was always done in the same order as the first three tasks. It should be noted, that, since there are six potential orderings of three tasks, but only 10 pairs in each condition, two orderings in each condition were used only once rather than twice.

4.4 Measures of Groupware Usability

We use five measures in this study: completion time, perceived effort, verbal efficiency, overall preference, and strategy use. These can be characterized as product, process, or satisfaction measures (see Table III).

Table III. Summary of Measures Used

Type of Measure	Measure Used
Product	Completion time
Process	Verbal efficiency, perception of effort, strategy use
Satisfaction	Overall preference

Table IV. Sex Pairings of Experimental Groups

Pairing	Overview Condition	Radar Condition
Male-Male	6 pairs	6 pairs
Female-Female	1 pair	1 pair
Female-Male	3 pairs	3 pairs

- (1) *Completion time* is a basic measure of product performance. It assumes that there is a relationship between the activities of collaboration and the speed at which a group can perform the task.
- (2) *Verbal efficiency* is a more direct measure of communication. Each of the tasks requires that participants convey a certain amount of information (e.g., the location of 10 sites to be welded), but the number of words that a group uses to convey that information can vary. Therefore, the number of words spoken (on the subject of site location, for example) indicates how efficiently participants are communicating that information.
- (3) *Perception of effort* is a subjective measure of the criterion of effort for the activities of collaboration. We recognize, however, that people will have difficulty differentiating between these activities, so the measure only collects overall information.
- (4) *Overall preference* is a broad satisfaction measure based on a comparison of the two systems. It assumes that there is a relationship between overall usability and preference: that participants will prefer a system that better supports the activities of collaboration.
- (5) *Strategy use* is a qualitative process measure that looks at how groups in the different conditions carried out the task. We assume that a more usable system will allow groups to choose more appropriate strategies for each task.

4.5 Participants

Participants were recruited from the student community at the University of Calgary, and were paid \$10. Participants ranged in age from 19 to 48 years (median age was 24 years). Forty people participated in the study, 30 men and 10 women. Although there were unequal numbers of female and male participants, sex pairings were equalized across the two conditions, as shown in Table IV.

Participants were assigned a partner for the study, either by choosing one themselves or by random assignment. Of the 20 pairs, 11 were familiar with one another (two or more interactions per week). This method of assignment was used in order to maximize the number of participants that could be recruited. Degree of familiarity was controlled in the study by equalizing the variable in the two conditions (six familiar pairs started in the overview condition, five in the radar).

Participants had limited prior experience with real-time groupware, although all participants used email systems and web browsers more than once per week. Eight of the 40 participants also had experience with multiplayer games, and these participants were equally distributed across the study conditions. None of the participants had previously seen the groupware system used in the study.

4.6 Procedure

Participants were first introduced to the system's functions. Pairs were then randomly assigned to either the radar or the overview condition, and the specifics of their miniature view were explained. These explanations used previously written instructions and were designed to avoid biasing people toward one view or the other. Participants were then allowed to practice with the system until they could each perform a basic set of simple tasks (selecting, dragging, scrolling, welding, and unwelding) to the experimenter's satisfaction.

Pairs then completed seven tasks with the pipeline system: three tasks with one version of the system (either radar or overview), and then three with the other version. For each task, a similar procedure was followed. First, the experimenter explained the task and the goal. Second, the pair completed a practice exercise for the task. Third, the pair carried out the task. Fourth, participants filled out a questionnaire relating to the task. After all tasks were completed, participants also filled out a final questionnaire relating to their preferences.

Pairs worked with both interfaces so that they could state their preference at the end of the session. Both interface order and task order were counterbalanced; the second three tasks were always done in the same order as the first three tasks.

Four types of data were collected:

- (1) Completion time for each task was recorded with a stopwatch.
- (2) Verbal communication was recorded on videotape, and parts were later transcribed.
- (3) Participants answered questions about perceived effort after each task. Questions used five-point scales with fixed endpoints (see Table VI). Questionnaires were completed by individuals rather than by groups.
- (4) Participants were asked their preference between the two systems after they had completed all tasks. Again, we collected these data as individual rather than pair responses.

Table V. Task Sequence for Radar and Overview Conditions

	Tasks 1-3 (radar)	Tasks 1-3 (overview)	Tasks 4-6 (radar)	Tasks 4-6 (overview)
Radar 1st	Pairs 1–10			
Overview 1st		Pairs 11–20	Pairs 11–20	Pairs 1–10

Table VI. Perceived-Effort Questionnaire

How difficult was it to complete this task?	difficult	● ● ●	easy
How much effort did this task require?	little effort	● ● ●	
How hard did you have to concentrate to do this task?	not hard	● ● ●	very hard
How difficult was it to discuss things during the task?	easy	● ● ●	difficult




Fig. 6. Experiment room setup.

4.7 Physical Setup

Participants worked at separate workstations, angled so that they could not see each others' screens, but so that they could see one another and talk easily. This situation simulates a distributed setting with a high-quality video link (although a true distributed situation would not provide nearly the same visual quality). The experimenter sat at a recording station at the back of the room. The actions of both participants were transmitted to a third computer that showed a composite of the workspace. This computer's screen and both voices were recorded on videotape. The layout of the experiment room is shown in Figure 6.

5. RESULTS

In the following sections, we report on the results of several analyses performed on product, process, and satisfaction measures. We first present the results of the between-participants hypothesis, based on data from each group's first workspace miniature and first three tasks only. Second, we present the results of the within-participants hypothesis, based on relative comparisons of similar tasks performed with both workspace miniatures.

Table VII. Summary of Completion Times (in minutes) for Tasks 1–3

Task	View	N	max	min	Mean	sd
Follow 1	Radar	10	4.73	2.05	3.21	0.84
	Overview	10	7.82	2.22	4.58	1.54
Copy 1	Radar	10	4.77	2.20	3.36	0.91
	Overview	10	4.52	1.70	3.12	0.90
Direct 1	Radar	10	4.20	2.38	3.19	0.63
	Overview	10	5.87	3.02	4.39	1.07

Third, we report participants' overall preferences, and last, we categorize and describe participants' strategy use.

In addition to these analyses, we checked for effects of two particular demographic variables: participant familiarity and experience with multi-user games. Neither of these variables showed statistically significant differences; that is, familiar groups were not significantly different from unfamiliar groups, and experienced gamers were not significantly different from those without multiuser game experience.

Finally, the data used below reflect the fact that two groups did not complete the second set of tasks, due to time restrictions, and that two groups did not have their conversation recorded due to technical problems. We do not believe that our analysis or conclusions are affected by these missing data.

5.1 Between-Participants Results

Our first hypothesis is that groups in the radar condition will complete the first three tasks more quickly, with greater verbal efficiency, and with less perceived effort. The hypothesis was tested by looking for effects of View, and the three measures taken were completion time, word counts, and subjective perception of effort. There were no main effects of View for any of these three measures. However, there were various interactions between View and Task, as described below.

5.1.1 Completion Time. Completion times were recorded for each task. Times for tasks 1–3 are summarized in Table VII, and shown in Figure 7 (error bars indicate standard deviation). Tasks took participants between about two and about eight minutes; for Follow and Direct tasks, the average completion time was less for the radar condition than for the overview condition.

We tested the effect of View using two-way analysis of variance (ANOVA). There was no main effect of View ($F = 4.29, p = 0.053$). There was an interaction between Task and View ($F = 7.77, p < 0.05$). Since the three kinds of tasks were quite different, as mentioned above, differences between task types were expected and were not analyzed. To explore the effect of View in the interaction, post hoc comparisons of radar and overview completion times were carried out for each task type. Our expectation was that the radar condition would have lower completion times;




Fig. 7. Mean completion times (in minutes) for tasks 1–3.

Table VIII. Comparisons of Completion Times for Tasks 1–3

Task Type	df	Tails	t (obtained)	<i>p</i>	r_{pb}^2
Follow 1	18	1	2.48	< 0.0167	0.255
Copy 1	18	2	-0.580	= 0.569	
Direct 1	18	1	3.05	< 0.0167	0.341

therefore, for the Follow and Direct tasks we used one-tailed t-tests. However, completion times for the Copy task were contrary to our expectation, so we could not use a one-tailed test. As a fallback for this task, we used a two-tailed test instead. A Bonferroni correction was employed to maintain alpha below 0.05; therefore, only those effects with $p < 0.0167$ were considered statistically significant. Of the three tasks, differences in Follow and Direct were significant. Results of the post hoc comparisons are summarized in Table VIII. The proportion of variance accounted for by View is indicated by the squared point-biserial correlation coefficient (r_{pb}^2). The coefficient indicates that only about one-quarter to one-third of the variance in the sample is accounted for by View.

5.1.2 Communication Efficiency. Verbal interaction in the first three tasks was recorded and transcribed, and the transcripts were coded by two assistants. Coding involved two steps: first, the transcript was divided into utterances (contiguous sequences of words dealing with a single intention); second, utterances were categorized into a number of classes. Communication efficiency was then measured by counting the number of words in particular categories. In the Follow and Direct tasks, the category of interest included any words that the group spoke in establishing where the welder or actor was to go next. For example, the category included words spoken by either participant that described locations, provided directions, or clarified location information. When the welder or actor arrived at the correct location, counting stopped: that is, the category did not include talk




Fig. 8. Mean verbal efficiency (in number of words) for tasks 1–3.

Table IX. Summary of Verbal Efficiency (in Number of Words) for Tasks

Task	View	N	max	min	Mean	sd
Follow 1	Radar	8	123	64	98.75	21.77
	Overview	10	348	103	221.43	77.09
Copy 1	Radar	8	224	0	129.08	86.55
	Overview	10	133	1	73.30	49.59
Direct 1	Radar	8	345	104	223.50	81.40
	Overview	10	427	138	280.97	98.09

about the location or orientation of particular pipe sections. For the Copy task, the category included any words that the group spoke in establishing which pipe section the copier was to select next. For example, the category included words spoken to indicate a piece, describe a piece, or clarify a description. Counting stopped once the copier selected the correct section.

Two assistants each coded half of the transcripts and counted the words in each category. On a test set of four transcripts, interrater agreement between the two coder's counts (using Pearson's r) was above 80% for all three tasks. Word counts are summarized in Table IX, and mean counts are illustrated in Figure 8 (error bars indicate standard deviation).

No main effect of View was found ($F = 3.60$, $p = 0.074$). Analysis of variance again showed an interaction between Task and View ($F = 9.75$, $p < 0.05$). To assess the effect of View on verbal efficiency, we compared radar and overview conditions for each task type. We used one-tailed t-tests for the Follow and Direct tasks; again, means for the Copy task did not meet our assumptions for using one-tailed tests, so a two-tailed test was used. The tests showed a significant difference only for the Follow task. A summary of the comparisons is shown in Table X.

5.1.3 Perceived Effort. Perception of effort was measured by a repeated questionnaire given after each task. The questionnaire looked at (1) overall difficulty, (2) effort required, (3) concentration required, and (4) difficulty

Table X. Comparisons of Verbal Efficiency for Tasks 1–3

Task Type	df	Tails	t (obtained)	<i>p</i>	r_{pb}^2
Follow 1	18	1	4.34	< 0.0167	0.541
Copy 1	18	2	-1.72	0.104	
Direct 1	18	1	1.32	0.101	

Table XI. Summary of Questionnaire Responses, Tasks 1–3

		Task: Follow 1 Question				Task: Copy 1 Question				Task: Direct 1 Question			
		1	2	3	4	1	2	3	4	1	2	3	4
Mean	Radar	1.65	2.05	2.55	1.70	1.90	2.30	2.40	2.15	1.90	2.40	2.50	2.20
	Overview	2.10	2.75	2.75	2.15	1.30	1.55	2.30	1.50	2.15	2.55	2.80	2.70
SD	Radar	0.75	0.94	0.94	0.86	0.91	1.03	1.05	1.09	0.79	0.99	0.95	0.95
	Overview	1.07	0.97	1.02	1.09	0.57	0.69	0.98	0.69	0.99	0.89	1.01	1.30

1. How difficult was it to complete this task?
2. How much effort did this task require?
3. How hard did you have to concentrate to do this task?
4. How difficult was it to discuss things during the task?

discussing the task. Questions used five-point scales with semantic anchors (see Table VI). Responses were translated to interval scores, using 1 to represent least effort and 5 to represent most effort.¹ Table IX summarizes mean responses for each question in each task, and Figure 8 illustrates the means. Note that lines connecting the points are intended only to visually differentiate the two conditions, not to imply connections between questions.

We used responses to question 2 (“How much effort did this task require?”) as an overall indicator of perceived effort. ANOVA showed no main effect of View ($F = 0.03$, $p = 0.86$), but showed an interaction between Task and View ($F = 7.34$, $p < 0.05$). To look for specific differences, we compared radar and overview responses for each question. Alpha of 0.05 was divided between the 12 tests; therefore, only results where $p < 0.0042$ were considered statistically significant. None of the comparisons showed significant differences. Comparisons are summarized in Table XII.

5.2 Within-Participants Results

Completion times and questionnaire results were also gathered for the second trio of tasks, those completed with the second workspace miniature. These measures allowed us to consider the issue of what happens when a

¹Converting categorical data to interval scores misrepresents the data to a certain degree; however, χ^2 analysis was not possible on the data due to low expected values.

Table XII. Comparisons of Perceived-Effort Questions

Question	Follow 1			Copy 1			Direct 1		
	df	Tails	<i>p</i>	df	Tails	<i>p</i>	df	Tails	<i>p</i>
1	38	1	0.656	38	2	0.0171	38	1	0.1910
2	38	1	0.0130	38	2	0.0101	38	1	0.3088
3	38	1	0.2619	38	2	0.7566	38	1	0.1686
4	38	1	0.0780	38	2	0.0299	38	1	0.0868

group moves from one view type to another. We assumed that all groups would perform better in the second set of tasks because of practice, but we wondered whether the improvement would be greater when going from the radar view to the overview, or when going from the overview to the radar view. We consider the differences between first and second trials for completion time and perceived effort. Verbal records were not transcribed for the second set of tasks, so differences in communication efficiency are not analyzed.

The within-participants data contain a potential confound of order. That is, within-participants differences cannot be completely explained in terms of View, due to the confound of training effects from the first interface. For example, groups who used the radar view first may have learned particular strategies that did not work with the overview, thus causing them to use more time in their second trials. Although we believe the effect of the confound to be small, we present these results primarily as supplementary findings.

5.2.1 Completion Time Differential. The time difference between a group's first and second attempts at a particular task indicates their improvement, and we expected that groups would be faster on their second attempt. Table XIII summarizes the differences between a group's first and second attempts at each task, and Figure 9 illustrates these differentials.

For groups that started with the overview and then moved to the radar view, the results were as we expected: for each task type, groups were faster in the second attempt (using the radar view). However, when groups started in the radar condition and then used the overview, only the Copy task was faster in the second attempt. The Follow and Direct tasks were both slower with the overview: Follow by about a minute, and Direct by about half a minute.

Analysis of variance indicates a main effect of View ($F = 8.52$, $p < 0.05$); observation of the data shows that the radar condition allowed a larger improvement in performance. No significant interactions were found.

5.2.2 Perceived-Effort Differential. A similar analysis was done with questionnaire responses. The difference between a participant's first and second responses to the questionnaire indicated whether they thought the second task was easier or harder than the first. Differentials were calcu-

Table XIII. Summary of Completion Times Differentials (in minutes)

Task	View Order	N	max	min	Mean	sd
Follow	Radar then overview	10	4.48	-0.60	0.68	1.57
	Overview then radar	9	0.17	-4.55	-1.44	1.42
Copy	Radar then overview	9	1.52	-1.37	-0.14	1.04
	Overview then radar	9	-0.12	-1.17	-0.51	0.36
Direct	Radar then overview	9	1.53	-0.43	0.47	0.57
	Overview then radar	9	-0.27	-1.98	-1.03	0.48

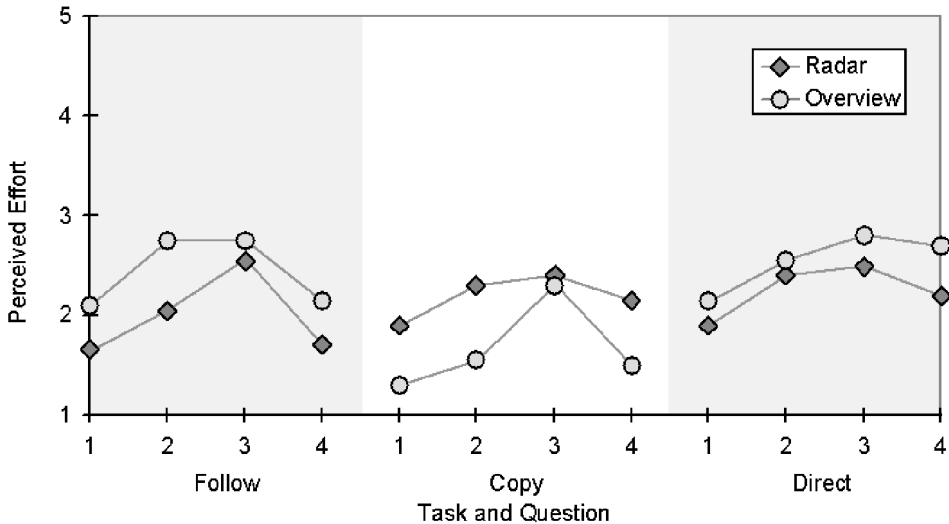


Fig. 9. Mean questionnaire responses for tasks 1-3.

lated by subtracting the first response from the second. Tables XIV–XVI summarize the mean differentials for each questionnaire question.

Figure 10 illustrates these differentials in perceived effort. In the figure, points below the zero line indicate that the second task was perceived to be easier than the first, and points above the line that the second task was perceived to be harder. Again, the lines in the figure are intended only to visually separate the two data sets.

We assumed that as groups became more experienced at each task, they would consider it to require less effort. This was the case when groups used the overview first and the radar second: they felt that the second task was easier. However, when groups used the radar and then the overview, they felt that the second task was more difficult than the first.

We again used responses to question 2 (“How much effort did this task require?”) as an overall indicator of perceived effort, and carried out analysis of variance on the differentials for question 2 from the first to the second task. ANOVA showed a main effect of View ($F = 27.83, p < 0.05$), and an interaction between Task and View ($F = 4.54, p < 0.05$).

Table XIV. Mean Questionnaire Differentials for the Follow Task

Question		1	2	3	4
Mean	Radar then overview	1.13	0.95	0.65	0.80
	Overview then radar	-0.83	-1.22	-0.94	-1.06
SD	Radar then overview	1.28	1.36	1.31	1.54
	Overview then radar	0.99	1.06	0.73	1.11

Table XV. Mean Questionnaire Differentials for the Copy Task

Question		1	2	3	4
Mean	Radar then overview	0.33	0.22	0.06	0.06
	Overview then radar	-0.06	-0.11	-0.56	-0.28
SD	Radar then overview	0.84	1.06	1.00	1.13
	Overview then radar	0.73	0.90	0.78	0.75

Table XVI. Mean Questionnaire Differentials for the Direct Task

Question		1	2	3	4
Mean	Radar then overview	1.44	1.06	1.17	0.89
	Overview then radar	-0.06	-0.11	-0.50	-0.22
SD	Radar then overview	1.04	1.11	0.71	1.08
	Overview then radar	1.30	1.23	1.42	1.40




Fig. 10. Mean changes in completion time from first to second attempts at a task.

5.3 Preference

After all tasks were completed and pairs had used both interfaces, participants were asked three questions about which system they preferred. The questions asked which system better supported collaborative work, which system was easier to use for group tasks, and which system the participant




Fig. 11. Perceived-effort differentials between first and second attempts at a task.

Table XVII. Number of Participants Preferring Each Interface

Which System...	Radar	Overview
(1) ...better supported your collaboration?	35	3
(2) ...was easier for group work?	38	0
(3) ...did you prefer overall?	38	0

Table XVIII. χ^2 Analysis of Preference Questions

Question	χ^2	df	<i>p</i>
1	26.95	1	$p < 0.0167$
2	38.00	1	$p < 0.0167$
3	38.00	1	$p < 0.0167$

preferred overall. Almost all of the participants who responded chose the radar view, as shown in Table XVII.

We analyzed these responses using one-way χ^2 tests, summarized in Table XVIII. Again, alpha was maintained at 0.05. Not surprisingly, the number of participants choosing the radar condition was significantly higher than the expected number for each question.

5.4 Strategy Use

We also looked at the strategies that groups used to carry out the tasks in the two conditions. In particular, we recorded the strategy used to indicate locations (for the Follow and Direct tasks) and to indicate pieces (for the Copy task). We identified strategies subjectively by watching the session videotapes. People used a wide variety of methods, both verbal and nonverbal, for indicating locations and pieces. The strategies are described in Table XIX.

Table XIX. Strategies Used for Directing and Indicating

Strategy	Used in	Description
Relative-to-you	Follow, Direct	Directions based on the other person's current location: e.g., "up and left from where you are"
Describe-location	Follow, Direct	A description of an object at the location: e.g., "the squiggly-looking thing"
Left-right-top-bottom	Follow, Direct	Rough coordinate system dividing the workspace into four blocks: e.g., "next one is in the top left corner"
Relative-to-previous	Follow, Direct	Directions based on a previously identified location: e.g., "near where we were for the last one"
Map-coordinates 3×3	Follow, Direct	Directions based on a 3-by-3 grid: e.g., "go to 1,2"
Pipe-tracing	Direct	Directions to follow a line of pipe: e.g., "follow this pipe along to the right, and then it goes up"
Follow-rectangle	Follow	One person tracks the other by following his or her view rectangle in the radar
Relative-to-us	Follow	Directions given when both participants are in the same place: e.g., "now down and a little to the left from here"
Move-piece-to-show	Follow	One person moves a pipe section to indicate a location through the radar or overview
1D-relative-and-wait	Follow	Directions to move up, down, left, or right, after which the person giving directions waits until success is established
Follow-my-cursor	Follow	One person follows the other's main view cursor
Describe-piece	Copy	A description of the next piece to be used: e.g., "it's an elbow section with a medium straight on the end"
Show-by-move	Copy	The piece is moved back and forth in the storehouse
Show-by-drag	Copy	The piece is dragged up to the construction area
Show-by-placing	Copy	The piece is moved to the construction area and placed

There were several differences in strategy use between the two conditions, differences that can be partly attributed to the information available in the two interfaces. Strategy use is summarized in Table XX. In general, groups in the overview used a wider range of strategies than groups in the radar condition. Strategies that we observed only in the overview condition include pipe-tracing (Direct task), 1D-relative-and-wait, follow-my-cursor, map-coordinates, and move-piece-to-show (Follow task). The only strategy seen solely in the radar condition was follow-rectangle (Follow task), which is understandable, since the overview did not provide a view rectangle to follow

5.5 Summary of Results

A variety of results were obtained, some showing improvement when there was additional awareness information, and some showing no difference between the two displays. With the between-participants data, no main

Table XX. Strategy Use in the Three Test Tasks

Strategy	Radar Condition	Overview Condition
Direct	relative-to-you left-right-top-bottom relative to previous	describe-location relative to previous relative-to-you pipe tracing
Follow	follow rectangle left-right-top-bottom relative to us relative to previous relative-to-you describe-location	left-right-top-bottom describe-location move piece to show relative to us relative to previous 1D and wait follow my cursor relative to you map coordinates (3 × 3)
Copy	show by drag show by move describe-piece	show by drop describe-piece show by placing

effects of View were found, although for each measure there was an interaction between View and Task. When using the radar view, groups finished the Follow and Direct tasks significantly faster, and used significantly fewer words in the Follow task. No differences were found in perceived effort for any of the tasks, and no differences were found on any measure for the Copy task. Analysis of the within-participants data supports the overall hypothesis (that awareness information improves usability); main effects of View were found for completion time and perceived-effort differentials, although these results contain a potential confound of training and order. In addition to the quantitative results, we also found that strategy use differed considerably between the conditions.

6. DISCUSSION

The two versions of the interface differed only in that the radar view provided visual indications of the other person's location, the location of his or her cursor, and the motion of objects that he or she moved. The statistically significant differences between these two very similar interfaces clearly suggests that the additional awareness information helped people complete some tasks more quickly and more efficiently. We interpret and explain these findings below. First, we consider two reasons why the additions to the radar view were successful: that they allow visual communication, and that they provide continuous feedback and feedthrough. Second, we examine the measures of perceived effort, and consider why the Copy task was not affected by the view type.

6.1 Visual versus Verbal Communication

The radar condition provided visual indication of the other person's location and activity by showing view rectangles and telepointers. This information

helped people complete the Follow and Direct tasks more quickly. One way that visual information aided the task was by allowing people to use strategies that were better suited to the task and therefore more effective.

6.1.1 Visual Information and Strategy in the Follow Task. In the Follow task, the joiner (the person with the map) had the job of communicating 10 successive workspace locations to the welder (who had no map). When groups used the overview system, the joiner had to convey this information verbally. Joiners used a wide variety of techniques for indicating locations, and were generally adept at choosing a technique that would best describe where the welder should go next. They often began with general directions (e.g., left-right-top-bottom or relative-to-previous strategies), and then gave more specific indication using the describe strategy. In many cases, however, the locations were not easy to indicate using any of the strategies. For example, when the next location was not obviously in a corner of the workspace, and not in an obvious direct line from the current position, then neither of left-right-top-bottom or relative-to-here were appropriate. In these situations, the joiner had to rely more heavily on describing the location and had to be more careful in planning and delivering her utterances. Often, the descriptions became fairly complicated:

J: The second weld is near the bottom in the middle section, there's two pieces of pipe, ok, there's two longer pieces of pipe, ok, there's, umm, right in the middle, right on top of the lowermost piece of pipe, in the middle there, there's two welds that need to be done.

W: Uh, ok...

The joiner's verbal instructions had to be interpreted by the welder, and this process took time. In addition, the joiner would sometimes have to provide more than one round of description before the welder found the correct location. In other cases, the problem was not incorrect interpretation, but incorrect direction:

J: Six is uh, down, to the right <J moves to the *left* side of the workspace>

W: <moves down and *right*>

J: Um, the very edge, there's one sticking down which is not welded

W: <looking for piece>

J: Uh, at the bottom?

J: See that?

W: No

J: Uh, oh, I mean, sorry, to the left, sorry

W: Oh, ok <moves left>

The radar view, in contrast, allowed people to use a much more effective strategy. The follow-rectangle strategy meant that the welder could find the right location simply by following the joiner's view rectangle. The visual indication of the joiner's location transformed the task from a series of complicated verbal exchanges to a relatively simple perceptual task of aligning rectangles on the screen. The follow-rectangle strategy provides specific and accurate information about where to go, regardless of where the next location is in the workspace. In addition, it allows the joiner to communicate simply by going about his or her job: they need not spend extra time thinking about how to best indicate the location.

The overview condition did in fact allow a limited kind of visual communication, but it was not as obvious as the follow-rectangle strategy and was not used very often. In the show-by-move strategy, the joiner would navigate to the next location, and then move a pipe section back and forth, knowing that each move would show up on the welder's overview. Although this strategy could provide a good indication of location, it could not be used consistently because there was not always a convenient section of pipe to move back and forth. In particular, where all the pipe sections in the area were connected into large structures, moving a structure would not provide an accurate indication of the joiner's location.

The transformation of the task from a verbal to a visual activity also explains why groups used significantly fewer words in the Follow task when they used the radar view. Groups using the follow-rectangle strategy had the necessary location information available in the radar, so they did not need to communicate locations verbally. In the audio record, the follow-rectangle strategy is characterized by few words, and almost none of the complicated and lengthy descriptions seen in the overview condition. However, the radar condition was never completely silent. In particular, joiners would often make general statements about the location of the next weld:

J: ok, we're going over to the left... that's getting welded

J: ok, now, way over here... ok, that needs to be welded

J: ok, and just over left, same height, weld this together...

Since these directions are too unspecific to fully indicate a location, the joiners must have been providing general directions but leaving the specifics up to the radar. In a few cases, when the joiner's directions became more specific, welders would remind them that specific directions were unnecessary, since the radar view provided the required information.

6.1.2 Visual Information and Strategy in the Direct Task. The Direct task also asked one person (the director) to communicate a series of successive workspace locations to the other person (the actor). Again, the director had a workspace map, and the actor had no map. In this task, the director was not allowed to move around in the workspace, so radar users could not employ the follow-rectangle strategy. However, even though the

director in both conditions had to indicate locations verbally, the information in the radar view allowed him or her to use more effective strategies.

As in the Follow task, workspace locations were not particularly easy to describe. Directors in the overview condition used several techniques to indicate locations, but still had some difficulty in indicating the right place to the actor, even though the actor could drop a piece to show his or her location:

D: Next, I need a small piece, from the bottom left

A: ok <gets piece>

D: and you want it right in the center, in that open space, there's a little pipe that sticks out

A: center in the open space...

D: in the top... you see one little pipe that sticks out, on the left?

A: ok, I'm here. <drops piece> Where do you want me to go?

D: ok, uh, up, to go to the top

A: the top right corner or the top left corner?

D: top left.

A: here? <drops piece>

D: yeah. Now go exactly right, from there. And you see, there's a T, with a pipe, straight?

A: There?

D: I can't see where you are.

Directors in the radar condition used many of the same strategies for indicating the next location (e.g., describe-location, left-right-top-bottom, relative-to-previous) as seen in the overview condition. However, when these strategies failed, location information in the radar view gave directors a fallback strategy that worked well even when locations were difficult to describe. Since directors could see exactly where the actor's view rectangle and telepointer were, they could provide relative directions (go up, go down, go left, go right) based on the actor's current location. Relative directions are simple to construct, and are much less prone to misinterpretation, such as

D: ok, move to the left, stop, stop. Move up, move straight up, move straight up, stop. Go a little bit to the left, stop, stop. Ok, now you see there are two T sections...

The relative-to-you strategy was not generally the first strategy chosen by a director, but it was often the one that he or she used when they ran into difficulty. In one session, the director started to describe the location, but after making a few attempts, resorted to relative directions:

D: Okay, number five. If you look at... there's some pipes to the... where, they're kind of... um...

D: Go down. <continues with relative directions>

The difference between descriptive or relative directions can also partly explain why the radar did *not* lead to fewer words spoken in the Direct task. Even though these two methods of giving directions differ greatly, nothing about giving relative directions implies that fewer words will be needed. For example, the first of the two utterances (D1) below might be harder to plan and to understand, but both utterances contain the same number of words. It may be that word counts are an insufficient measure of verbal efficiency, and that other metrics like utterance length or vocabulary size may have been more appropriate.

D1: ok, near the very bottom you'll notice that there's a vertical line right in the middle in the bottom of the pipeline, ok there is a T, a T, under that corner piece...

D2: ok, move to the left, stop, stop. Move up, move straight up, move straight up, stop. Go a little bit to the left, stop, stop. Ok, now you see there are two T sections.

In summary, the location information presented in the radar view allowed people to communicate required information visually in the Follow and Direct tasks. The visual information allowed different strategies for carrying out the tasks, and allowed simplification of verbal utterances. In the Follow task, the view rectangle was of primary importance in helping people complete the task more quickly; in the Direct task, both the view rectangle and the telepointer were important. This difference can be ascribed to the fact that pairs in the Follow task could use their main view to negotiate local directions, so the radar view was most useful in aligning views. In contrast, directors in the Direct task could only gather information about their partner through the miniature, so they had to provide both large-scale and small-scale directions using the miniature.

6.2 Continuous Feedback and Feedthrough

The radar view provided continuous feedback about location and piece position, feedback that allowed groups to complete the Follow and Direct tasks more quickly. In particular, this feedback gave people visual evidence of understanding [Brennan 1990], which was more effective and less error-prone than verbal evidence.

In the Direct task, the director guides the actor's movement by giving her an instruction. With each instruction, the director requires evidence that he has succeeded in conveying the correct meaning to the actor, and that the actor has successfully moved where she is supposed to go. In addition, the director often cannot give the next instruction until he knows that the actor has successfully completed the current one. The information differences between the radar view and the overview provided directors with

different kinds of evidence, and afforded different means for establishing that instructions had been understood and carried out.

The overview lets the actor give evidence in two ways: verbal acknowledgment (e.g., “ok, I’m there”) or the “here-I-am” strategy of dropping an object to indicate location (e.g., “ok, can you see my piece?”). In both of these methods, the evidence is given at the end of an action: that is, the director gives the instruction, and the actor carries it out to the best of her ability before acknowledging. The problem with this form of interaction is that the director may give poor descriptions, possibly causing the actor to go the wrong way. Providing evidence only at the end of the action means that time is wasted when the actor makes a mistake:

D:... go up to that part that’s jetting across the middle...

A: <moves>

A: <drops piece> this part right here?

D: Uh, on the left side actually, on the left side...

In addition, both the verbal and the “here-I-am” methods of acknowledgment have other drawbacks. If the actor believes she has followed the instruction correctly, but really has not, she will mislead the director with her acknowledgment. The director has little chance to detect the error, and so may continue, piling error upon error. The “here-I-am” strategy at least gives the director concrete information about the actor’s location, but this information can be out-of-date. Actors would often drop objects, then pick them up and keep moving. The director, however, saw only the out-of-date picture of the dropped piece. If he assumed that the location of the piece was also the location of the actor, errors could ensue.

The awareness information in the radar provided different kinds of evidence. Verbal acknowledgment was still possible, but the radar also showed up-to-the-moment object movement and viewport location. In the Direct task, these representations could be used as immediate visual evidence of the actor’s understanding and intentions. If the actor started moving the wrong way, the director would see the misunderstanding immediately:

D: ok, just above where you were working before...

A: <begins moving>

D: oh, not too far... yep, right... nope, up, up, up, higher, yeah, right there

The availability of continuous evidence also made it possible for people to give continuous instructions. This is a strategy with far fewer verbal turns, and where the actor acknowledges implicitly through actions. Clark [1996, p. 326] summarizes the difference between verbal and visual acknowledgment for on-going “installment” utterances like instructions:

in installment utterances, speakers seek acknowledgments of understanding (e.g., “yeah”) after each installment and formulate the next installment contin-

gent on that acknowledgment. With visual evidence, [the speaker] gets confirmation or disconfirmation while he is producing the current installment.

In summary, evidence of understanding and action in the radar was accurate, easy to get, and timely. The director was able to determine more quickly whether the instruction was going to succeed, and could reduce the cost of errors.

6.3 Perceived Effort

Measures of perceived effort in the between-participants analysis showed no differences between the two conditions for any task. This runs contrary to both our expectations and our observations. We observed groups having more difficulty discussing the task, and making more errors, when they used the overview. It is possible that the questionnaire was a poor measure of effort. The main problem was that the questionnaire asked people to make absolute judgments of effort; however, people are in general better at making relative comparisons (i.e., “which required more effort, A or B?”) than they are at absolute judgments (i.e., “how much effort did it require?”). In participants’ first set of responses (those used for the between-participants analysis), people had nothing to compare their experience to, and may have been unable to accurately indicate their effort on the scales given.

In the within-participants analysis, however, we did find a main effect of View. Once participants had seen both interfaces, questionnaire responses showed greater differences (see Figure 11). Therefore, this data did include an implicit relative component; in people’s second set of responses, they had their first experience to compare to. From these experiences, we suggest that subjective responses about issues like effort and satisfaction will be more sensitive to differences between conditions if participants are allowed to compare rather than simply rate their experience.

6.4 Explaining the Copy Task

In the Copy task, the two participants built two identical structures from two stockpiles. The leader had a paper picture of what was to be built, and had to indicate each successive pipe section to the copier, who had no picture. We found a main effect of View in the within-participants analysis, but the between-participants analysis for the Copy task showed no effects of View on any measure. There are several possible reasons why the additional awareness information did not improve performance or efficiency, and the most important of these again concerns strategy. The strategy that a group chose for the Copy task had a large impact on their completion time and their verbal efficiency, regardless of which interface they used. Participants typically used one of two strategies to indicate the next piece to their partners: they could describe the piece verbally (describe-piece), or they could show it to them through the radar or overview (show-by-drag or show-by-drop). Describing pieces was certainly the word-

ier strategy and was also slower, so strategy choice may have overshadowed any effect of the workspace awareness information.

In addition, it turned out that there were efficient strategies available in both the radar and overview conditions. The show-by-drag and show-by-drop strategies provide almost the same information to the person doing the copying. However, since show-by-drop is a less obvious strategy than show-by-drag, we had expected describe strategies to be more prevalent in the overview condition. However, this was not the case. Even though the radar view allowed people to point out pieces quite easily, the video record suggests that more groups used the describe-piece strategy in the radar condition than in the overview condition. In a few cases, choosing to describe rather than show pieces seemed to be the result of inexperience: during one session, the leader said “oh right—I keep forgetting that we can both see the same radar view,” whereupon she switched from a describe to a show strategy.

The combination of an equivalent strategy in the overview condition and a greater use of description in the radar condition account for the lack of speed or efficiency differences between the two conditions for the Copy task. The within-participants analysis did show an effect of View, and informal analysis of the session records suggests that people were more likely to choose a “show” strategy in their second Copy task, which reduces the overshadowing effect of strategy choice. It is also noteworthy that even when the addition of awareness information did not improve the task, neither did it significantly impair people’s performance.

7. LESSONS FOR GROUPWARE DESIGNERS

There are several lessons that groupware designers can take from this study. First, the findings reaffirm the value of workspace miniatures, as suggested by our previous qualitative study [Gutwin et al. 1996]. In the present experiment, we regularly observed people using their miniature view to orient themselves in the workspace, to navigate, to keep track of the current global state of the activity, and to carry out individual work that did not fit inside the main view. These and our earlier observations clearly suggest that many kinds of shared-workspace groupware systems will benefit from a workspace miniature, even if it is not used as a vehicle for awareness information. Second, the main finding of the study is that adding workspace awareness information to the miniature—visual indications of viewport location, cursor movement, and object movement—can significantly improve speed, efficiency, and satisfaction. These awareness components should be included in shared-workspace applications.

The study provides evidence that workspace awareness information can be useful in certain situations. We now address how these results generalize to the real world of collaborative work; we consider the similarity of both the participants and the tasks to the real-world users and tasks where awareness techniques may be used.

7.1 Participants

The participants in the study were university students, whereas the target population is a fairly broad group—anyone who will use shared-workspace groupware in a work situation. Our participants do share several characteristics with more realistic users: for example, they have considerable basic experience with small-group collaboration, and they are experienced with mouse-and-windows software. The main difference, however, is in the expertise that real users may have with the domain and the particular workspaces used for collaboration. The more that people work with particular tasks and particular artifacts, the better they will become at activities such as following, copying, and directing—and the less they will need an explicit presentation of workspace awareness information. For example, if the Follow task was always carried out in the same workspace, people would begin to find names for locations that were difficult to describe at first [Krauss and Fussell 1990], and it would become simple to meet at a particular location simply by naming it to the other person (e.g., “meet me at the upside-down staircase”). This indicates that workspace awareness information will be of more value to users who are not experts in the domain or to users who are not familiar with the layout of the artifacts in the workspace. This will certainly occur in many settings: for example, CAD draftspeople may have to make revisions to an existing drawing that is not familiar to them, and architects and builders may collaborate over many different blueprints over the course of a workday. As a general rule, greater uncertainty in a person’s knowledge about the workspace or the task will imply more value in explicit presentation of workspace awareness information.

7.2 Tasks

The experimental scenario and the test tasks were highly controlled, and so will not be exactly the same as any in the real world. However, the tasks were chosen to involve basic kinds of joint activities that we had seen in real-world collaboration, and so do have applicability to other situations. The tasks involve actions like moving to where someone else is, indicating an object to another person, or telling someone how to get to a specific location. We believe that these activities are extremely common in shared-workspace collaboration: for example, it would not be unusual to overhear statements like “come with me, let’s go look at that” (follow), “what is *this* thing” (indicate), or “it’s over there—further, further, right there” (direct) in a wide variety of collaborative settings. As further illustration, the following examples show how realistic tasks might involve these activities:

—An engineer at a main office is coordinating the installation of new machinery at a production site. One task in this job is to collaborate with a technician at the site to check the accuracy of the plant’s layout drawings (to ensure there will be enough space for the machinery). The engineer and technician use a layout drawing as their shared workspace, and the engineer asks the technician to check the true location of

particular components. For each measurement that the engineer wants checked, she must succeed in having the technician join her at the appropriate part of the drawing (similar to the Follow task) before she can make her request.

—A senior and a junior designer are working together on a VLSI circuit schematic. The senior person knows the diagram and the job well, but the junior designer does not, and so the senior is assigning the junior small tasks during the session. The senior designer regularly tells the junior to go to a particular part of the diagram (directing) to do these small tasks, and the junior partner regularly asks for assistance in determining the placement and function of particular components (indicating).

Given that the activities do occur in real-world collaboration, we must still consider how often they occur, and whether workspace awareness information will actually assist people in carrying them out. In tasks where information about locations and activities is needed, and where that information is difficult to provide verbally, the radar view will have a positive effect. However, the size of the effect on real-world tasks depends upon what portion of the task can benefit from visual information and continuous feedback. In Follow and Direct, the radar condition was faster by about 25%, a substantial margin. However, realistic tasks will include a mix of different activities, some that will benefit from the awareness information, and some that will not. Although the information will still be useful for part of the task, differences will be harder to quantify. In addition, expertise and familiarity with the workspace may reduce the need for explicit awareness information. The positive effects of workspace awareness information, therefore, can only be determined by considering these issues in terms of a specific work context. Although the effects may be less than those shown in this study, we believe that the information will still improve usability in many systems.

The value of explicit workspace awareness information, such as that given by the radar view, can only be determined by understanding both the awareness requirements of the task and the strategies that users have evolved to maintain that awareness. The Copy task provides a good example of how insufficient understanding of the task and the users can lead to mixed results in an evaluation. Although we knew that participants had to indicate pieces to one another, we did not realize that there were adequate ways to do this using the ordinary overview (the “show by dropping” strategy). This is a case where the awareness requirements are already met fairly well by existing facilities, and where the addition of new information did not make an obvious difference. We believe a similar situation would have emerged in the Follow and Direct tasks if workspace locations had had well-known names; this would have allowed people to meet their awareness needs through verbal communication alone. Even in these cases we believe that workspace awareness information will not detract from the task, and will add some value, even if only as redundant

feedback that helps confirm what people are receiving through other channels.

In summary, we believe that workspace awareness information, at worst, will not detract from a collaborative task; and at best, it can significantly improve a group's performance and satisfaction with the system.

8. CONCLUSION

In this research, we examined the hypothesis that interface support for workspace awareness can improve groupware usability. We carried out an experiment to look at the effects of showing viewports, cursors, and object motion in a workspace miniature. For tasks that use information about location and activity, and where constructing verbal descriptions is difficult, the workspace awareness information in the radar can reduce completion time, improve communicative efficiency, and increase satisfaction. The improvements in speed and verbal efficiency can be explained in terms of visual communication and continuous feedback. Visual information about location allows groups to use more effective and more robust strategies for carrying out the Follow and Direct tasks. Continuous feedback on people's movement through the workspace allows people to recognize and correct navigational errors quickly. The study adds quantitative evidence to the qualitative findings of our prior study [Gutwin et al. 1996], and begins to put intuitions about awareness onto an empirical footing.

Our further research in this area will move in two directions. First, we will continue work on quantitative evaluations of groupware usability. Some of the questions that we were unable to explore in this experiment include the effects of awareness support in other kinds of tasks such as organization or creation, and how well the radar view works when there are more than two people in the group. Second, we want to look more closely at the links between shared workspaces, communication, and collaborative interaction. Work in this direction will look more carefully at naturalistic situations and use methods like conversation analysis and interaction analysis (e.g., Suchman and Trigg [1991]). One possible analysis in this direction will be to look more closely at errors in coordination and communication, something that we observed in this study but did not measure. We believe that the connection between communication and the environment can tell us a great deal about groupware usability and about the information requirements of the next generation of groupware systems.

SOFTWARE

GroupKit, the groupware toolkit used to build the pipeline system, is freely available from www.cpsc.ucalgary.ca/projects/grouplab. The pipeline system itself can be obtained by contacting the authors.

ACKNOWLEDGMENTS

Thanks to Mark Roseman, Jase Chugh, Krista McIntosh, and Jeff Caird for discussions about and assistance with the system, the study, and the

analysis. Thanks to the anonymous referees for their suggestions and recommendations.

REFERENCES

- ANDERSON, A., O'MALLEY, C., DOHERTY-SNEDDON, G., LANGTON, S., NEWLANDS, A., MULLIN, J., FLEMING, A., AND VAN DER VELDEN, J. 1997. The impact of VMC on collaborative problem solving: An analysis of task performance. In *Video-Mediated Communication*, K. E. Finn, A. J. Sellen, and S. B. Wilbur Lawrence Erlbaum Assoc. Inc., Hillsdale, NJ, 133–155.
- BAEKER, R. 1993. *Readings in Groupware and Computer-Supported Cooperative Work*. Morgan Kaufmann, San Mateo, CA.
- BEDERSON, B. B. AND HOLLAN, J. D. 1994. Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology (UIST '94, Marina del Rey, CA, Nov. 2–4)*, P. Szekely, Ed. ACM Press, New York, NY, 17–26.
- BRENNAN, S. 1990. Seeking and providing evidence for mutual understanding. Ph.D. Dissertation. Stanford University, Stanford, CA.
- CHAPANIS, A. 1975. Interactive human communication. *Sci. Am.* 232, 3 (Mar.), 36–42.
- CLARK, H. 1996. *Using Language*. Cambridge University Press, New York, NY.
- DOERRY, E. 1997. An empirical comparison of copresent and technologically-mediated interaction based on communicative breakdown. Ph.D. Dissertation. Department of Computer and Information Science, University of Oregon, OR.
- DOURISH, P. AND BELLOTTI, V. 1992. Awareness and coordination in a shared workspace. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4)*, M. Mantel and R. Baecker, Eds. ACM Press, New York, NY, 107–114.
- EGIDO, C. 1990. Teleconferencing as a technology to support cooperative work: Its possibilities and limitations. In *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, J. Galegher, R. E. Kraut, and C. Egido, Eds. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 351–372.
- GAVER, W. 1991. Sound support for collaboration. In *Proceedings of the 2nd European Conference on Computer-Supported Cooperative Work (ECSCW '91, Sept. 25–27)*, L. Bannon, M. Robinson, and K. Schmidt, Eds. Kluwer B.V., Deventer, The Netherlands, 293–308.
- GAVER, W. W., SMITH, R. B., AND O'SHEA, T. 1991. Effective sounds in complex systems: The ARKOLA simulation. In *Proceedings of the Conference on Human Factors in Computing Systems: Reaching through Technology (CHI '91, New Orleans, LA, Apr. 27–May 2)*, S. P. Robertson, G. M. Olson, and J. S. Olson, Eds. ACM Press, New York, NY, 85–90.
- GREENBERG, S., Ed. 1991. *Computer-Supported Cooperative Work and Groupware*. Academic Press Computers and People Series. Academic Press Ltd., London, UK.
- GRUDIN, J. 1990. Groupware and cooperative work: Problems and prospects. In *The Art of Human Computer Interface Design*, B. Laurel, Ed. Addison-Wesley, Reading, MA, 171–185.
- GUTWIN, C. 1997. Workspace awareness in real-time distributed groupware. Ph.D. Dissertation. University of Calgary, Calgary, Canada. Available via <http://www.cs.usask.ca/faculty/gutwin/publications>.
- GUTWIN, C. AND GREENBERG, S. 1996. Workspace awareness for groupware. In *Proceedings of the CHI '96 Conference Companion on Human Factors in Computing Systems: Common Ground (CHI '96, Vancouver, British Columbia, Canada, Apr. 13–18)*, M. J. Tauber, Ed. ACM Press, New York, NY, 208–209.
- GUTWIN, C. AND GREENBERG, S. 1998. Effects of awareness support on groupware usability. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '98, Los Angeles, CA, Apr. 18–23)*, M. E. Atwood, C.-M. Karat, A. Lund, J. Coutaz, and J. Karat, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 511–518.
- GUTWIN, C. AND GREENBERG, S. 1999. A framework of awareness for small groups in shared-workspace groupware. Tech. Rep. 99-3. Department of Computer Science, University of Saskatchewan, Canada.

- GUTWIN, C., ROSEMAN, M., AND GREENBERG, S. 1996. A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, G. Olson, J. Olson, and M. S. Ackerman, Eds. ACM Press, New York, NY, 258–267.
- HEATH, C., LUFF, P., AND SELLEN, A. 1997. Reconfiguring media space: Supporting collaborative work. In *Video-Mediated Communication*, K. E. Finn, A. J. Sellen, and S. B. Wilbur Lawrence Erlbaum Assoc. Inc., Hillsdale, NJ, 323–347.
- ISHII, H., KOBAYASHI, M., AND GRUDIN, J. 1992. Integration of inter-personal space and shared workspace: ClearBoard design and experiments. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4)*, M. Mantel and R. Baecker, Eds. ACM Press, New York, NY, 33–42.
- KRAUSS, R. M. AND FUSSELL, S. R. 1990. Mutual knowledge and communicative effectiveness. In *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, J. Galegher, R. E. Kraut, and C. Egido, Eds. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 111–145.
- MCDANIEL, S. E. AND BRINCK, T. 1997. Awareness in collaborative systems: A CHI 97 workshop report. *SIGCHI Bull.* 29, 4 (Oct.), 68–71.
- MONK, A., MCCARTHY, J., WATTS, L., AND DALY-JONES, O. 1996. Measures of process. In *CSCW Requirements and Evaluation*, P. Thomas, Ed. Springer-Verlag, Vienna, Austria, 125–139.
- NIELSEN, J. 1993. *Usability Engineering*. Academic Press Prof., Inc., San Diego, CA.
- O'CONNAILL, B. AND WHITTAKER, S. 1997. Characterizing, predicting, and measuring video-mediated communication: A conversational approach. In *Video-Mediated Communication*, K. E. Finn, A. J. Sellen, and S. B. Wilbur Lawrence Erlbaum Assoc. Inc., Hillsdale, NJ, 107–131.
- OLSON, J. S., OLSON, G. M., AND MEADER, D. K. 1995. What mix of video and audio is useful for small groups doing remote real-time design work?. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '95, Denver, CO, May 7–11)*, I. R. Katz, R. Mack, L. Marks, M. B. Rosson, and J. Nielsen, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 362–368.
- OLSON, J. S., OLSON, G. M., STORRØSTEN, M., AND CARTER, M. 1992. How a group-editor changes the character of a design meeting as well as its outcome. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4)*, M. Mantel and R. Baecker, Eds. ACM Press, New York, NY, 91–98.
- OLSON, J., OLSON, G., AND MEADER, D. 1997. Face-toface group work compared to remote group work with and without video. In *Video-Mediated Communication*, K. E. Finn, A. J. Sellen, and S. B. Wilbur Lawrence Erlbaum Assoc. Inc., Hillsdale, NJ, 157–172.
- ROSEMAN, M. AND GREENBERG, S. 1996. Building real-time groupware with GroupKit, a groupware toolkit. *ACM Trans. Comput. Hum. Interact.* 3, 1, 66–106.
- SOHLENKAMP, M. AND CHWELOS, G. 1994. Integrating communication, cooperation, and awareness: the DIVA virtual office environment. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '94, Chapel Hill, NC, Oct. 22–26)*, J. B. Smith, F. D. Smith, and T. W. Malone, Eds. ACM Press, New York, NY, 331–343.
- SHACKEL, B. 1990. Human factors and usability. In *Human-Computer Interaction*, J. Preece, Ed. Prentice Hall Press, Upper Saddle River, NJ, 27–41.
- SMITH, R. B. 1992. What you see is what I think you see. *SIGCUE Outlook* 21, 3 (Spring 1992), 18–23.
- 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 (ECCSCW '89, Gatwick, UK)*
- SEGAL, L. 1995. Designing team workstations: The choreography of teamwork. In *Local Applications of the Ecological Approach to Human-Machine Systems*, P. Hancock, J. Flach, J. Caird, and K. Vicente, Eds. Lawrence Erlbaum Assoc. Inc., Hillsdale, NJ, 392–415.
- STEFIK, M., BOBROW, D. G., FOSTER, G., LANNING, S., AND TATAR, D. 1987. WYSIWIS revised: Early experiences with multiuser interfaces. *ACM Trans. Inf. Syst.* 5, 2 (Apr.), 147–167.

- SUCHMAN, L. A. AND TRIGG, R. H. 1992. Understanding practice: Video as a medium for reflection and design. In *Design at Work: Cooperative Design of Computer Systems*, J. Greenbaum and M. Kyng, Eds. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 65–90.
- TANG, J. C. 1991. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud.* 34, 2 (Feb. 1991), 143–160.
- TATAR, D. G., FOSTER, G., AND BOBROW, D. G. 1991. Design for conversation: Lessons from Cognoter. *Int. J. Man-Mach. Stud.* 34, 2 (Feb. 1991), 185–209.
- WHITTAKER, S. 1996. Talking to strangers: An evaluation of the factors affecting electronic collaboration. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, G. Olson, J. Olson, and M. S. Ackerman, Eds. ACM Press, New York, NY, 409–418.

Received: October 1998; revised: July 1999; accepted: November 1999