A Usability Study of Awareness Widgets in a Shared Workspace Groupware System

Carl Gutwin, Mark Roseman and Saul Greenberg

Research Report No. 96/585/05

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

March, 1996

Gutwin, C., Roseman, M. and Greenberg, S. (1996).
A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. Report 1996-585-05, Dept Computer Science, University of Calgary, Calgary, AB, Canada. March

A Usability Study of Awareness Widgets in a Shared Workspace Groupware System

Carl Gutwin, Mark Roseman, and Saul Greenberg
Department of Computer Science, University of Calgary
2500 University Dr. NW, Calgary, AB, Canada T2N 1N4
Tel: +1 403 220-6015
E-mail: [gutwin, roseman, saul]@cpsc.ucalgary.ca

ABSTRACT
Workspace awareness is knowledge about others' interaction with a shared workspace. Groupware systems provide only limited information about other participants, often compromising workspace awareness. This paper describes a usability study of several widgets designed to help maintain awareness in a groupware workspace. These widgets include a miniature view, a radar view, a multi-user scrollbar, and a "what you see is what I do" view. The study examined the widgets' information content, how easily people could interpret them, and whether they were useful or distracting. Experimenter observations, subject questionnaires, and interviews indicate that the miniature and radar displays are useful and valuable for tasks involving spatial manipulation of artifacts.

Keywords
Workspace awareness, shared workspaces, usability study

INTRODUCTION

People maintain an ongoing awareness of others in physical workspaces like whiteboards and tabletops, and they do this using everyday perceptual abilities. For example, we can glance over at another person to see where they are working, or we might hear a particular tool that indicates what they are doing. In the virtual workspaces provided by real-time distributed groupware, these abilities are greatly reduced. Groupware systems reduce a person's visual field to the limited area of a computer screen, remove characteristic motions and sounds from actions, and complicate auditory and visual communication. The situation is made worse by groupware techniques like relaxed-WYSIWIS view linking, that can hide people's visible actions from one another.

As a result of these changes, people receive only a fraction of the information about others that they would in a face-to-face setting, and it becomes much more difficult to maintain awareness. One kind of awareness that is often compromised in the move to a groupware system is *workspace awareness*: the up-to-the minute knowledge a person holds about another's interaction with the workspace [13]. This includes knowledge about who is in the workspace, where they are working, what they are doing, and what they intend to do next. Workspace awareness reduces the effort needed to coordinate tasks and resources, helps people move between individual and shared activities, provides a context in which to interpret other's utterances, and allows anticipation of others' actions.

Several CSCW projects have considered support for various kinds of awareness, including workspace awareness (e.g. [2,6,7]). One approach is to augment the groupware interface with new components—widgets—that show some of the missing information about other collaborators. For example, telepointers are commonly added to WYSIWIS workspaces to show people's location and focus of attention. In relaxed-WYSIWIS systems, however, techniques like telepointers do not always work since participants may be looking at different areas of the workspace. Awareness widgets that have been created specifically for relaxed-WYSIWIS workspaces include:

- radar views (e.g. [14, 1, 18]);
- group versions of existing widgets, such as multi-user scrollbars [1];
- graphical activity indicators (e.g. [22,2])
- auditory cues (e.g. [4,10,2]);

However, only a few techniques have been tested, and then usually as part of an overall system evaluation. It is still unclear whether add-on widgets improve the effectiveness or usability of a relaxed-WYSIWIS groupware system.

We have conducted a study that examines several awareness widgets added to a real-time distributed groupware system. In this study, we look at three issues that underlie the usability of this approach.

1. Does the augmented system present the right type and amount of information to the user?
2. Can this information be easily interpreted and applied?
3. Does the additional information intrude on individual work, by using up screen space or by distracting people from their tasks?

Our goals were to answer these three questions, and to determine whether users found the widgets to be effective and valuable. The following sections describe related research on workspace awareness, the methodology of the study and the widgets used, and the results of the investigation. We then discuss possible explanations of our

results, their impact on designers of real-time groupware, and limits to the generality of our findings.

RELATED WORK ON WORKSPACE AWARENESS

The idea of helping people maintain a sense of awareness about distributed collaborators has been around since Englebart [9] used telepointers and video images in his shared workspace systems. More recently, several types of awareness in group work have been identified and studied in CSCW:

- *Informal awareness* of a work community is the general sense of who is around and what they are doing—the knowledge that people have when they work together in the same office (e.g. [6,3])
- *Social awareness* is information maintained about another person in a social or conversational context: their emotional state, their level of interest, or whether they are paying attention (e.g. [17]).
- *Structural awareness* involves knowledge about people's roles and responsibilities within a group, their positions on an issue, their status, and the state of group processes (e.g. [5]).

A fourth kind is *workspace awareness*, which is different from the other forms because of the workspace's integral role in affording and mediating interaction. Workspace awareness brings together several strands of CSCW research on view linking, group interface design, coupling in collaboration, and interaction in shared workspaces.

Issues of workspace awareness have surfaced in many groupware projects. Work on early WYSIWIS systems (e.g. [22]) found that people needed up-to-date information about others (such as who had just made a contribution) to collaborate effectively. In addition, studies of face-to-face collaboration showed the prevalence of workspace actions such as gesturing, listing, and drawing [21] and the importance of being able to see those actions. Ellis and colleagues [8], in their work building interfaces for group writing, recognized that differences between face-to-face work and groupware could imply a tradeoff between being well-informed about others' activities and being distracted from individual tasks. The distraction problem arises because groupware participants "are generally not as aware of others' contexts and can less easily interpret sudden display changes resulting from others' actions" (p. 49). They conclude that "a good group interface should depict overall group activity and at the same time not be overly distracting."

With relaxed-WYSIWIS groupware, lack of awareness about "others' contexts" became a bigger problem, since increasing individual control reduces the group focus inherent in strict-WYSIWIS systems. However, researchers realized that even when people work in a loosely-coupled mode, they remain connected to each other through awareness: "often, people are merely aware of each other—aware of others' presence, perhaps their activities and progress" [11] (p. 293). Gaver recognized the importance of awareness in helping people "shift from working alone to working together, even when joined on a shared task." As he says, "building systems that support these transitions is important, if difficult."

Dourish and Bellotti [7] apply these ideas more specifically to shared workspaces, and define awareness as "an understanding of the activities of others, which provides a context for your own activity" (p. 107). They also argue that awareness information should be passively collected and distributed rather than explicitly provided by participants, and it should be "presented in the same shared work space as the object of collaboration."

Drawing on these experiences, we have constructed a framework of workspace awareness that expands and adds precision to the concept [13]. The framework divides workspace awareness into several elements, each a kind of information that people maintain about others in a shared workspace. For example, people may keep track of who is present in the workspace, their location, their current activities, and their intentions. In this study, we concentrate primarily on support for awareness of location and activity.

METHODOLOGY

This section describes the subjects who used the widgets, the organization and setup of the study, the groupware application, and the awareness widgets that we tested.

Subjects

Nine pairs of computer science undergraduate and graduate students participated in the study as paid volunteers. The first pair acted as pilot subjects to test the groupware system and the experimental methods. Four of the eight remaining pairs had experience working with each other in class groups during the previous semester. All of the subjects were familiar with the workstation, optical mouse, and window system used in the study. Most were familiar with the idea of groupware, but none had experience with real-time distributed groupware or the particular system used in the study.

Physical setup

Participants worked at separate Sun workstations with 19-inch colour monitors. The computers were in the same room, but separated by a divider so that participants could not see either the other person or the other workstation. However, participants could talk normally across the divider. Two experimenters were present, and sat a few metres behind the participants, making notes and providing assistance when needed. A video camera recorded one workstation's screen and the subjects' conversation.

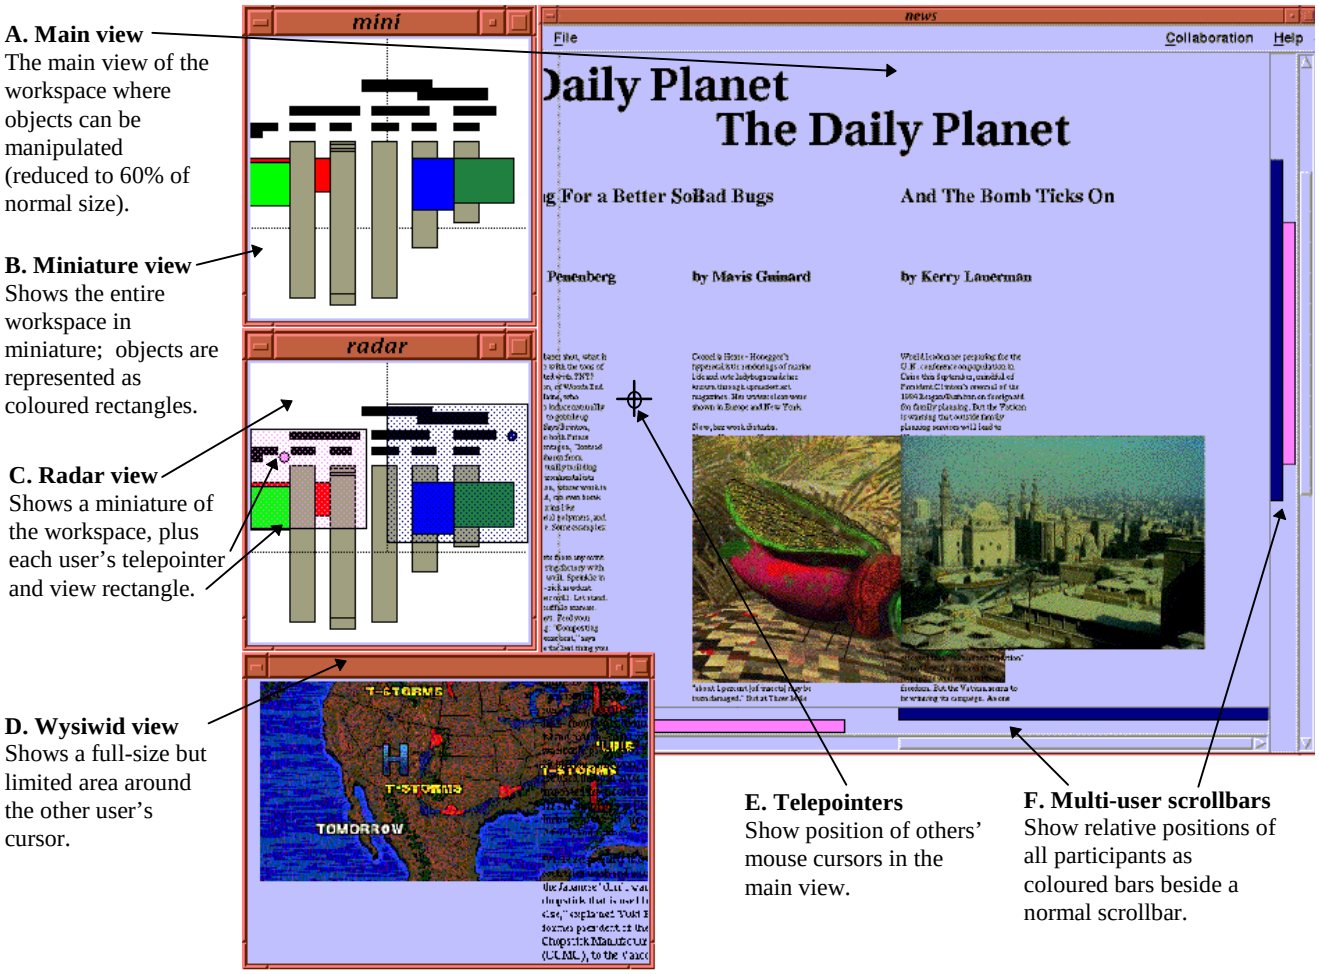

Figure 1. Awareness widgets and main view

Groupware system

A shared-workspace page layout application was built for the study, using the GroupKit toolkit [16]. The system's workspace simulated a two-page newspaper spread; the artifacts in the workspace included headlines, columns of text, pictures, and banners (see Figure 1A). Participants could move artifacts by dragging them, and could cut or join columns of text using the mouse.

Each person's main view occupied most of their screen, but showed only about one-third of the whole workspace. Participants could move their main view using scrollbars, and were able to position their views independently. In addition to the main view, various configurations of the system were augmented by awareness widgets, shown in Figure 1 and described below.

Widgets

In all configurations of the system, we included two techniques as basic support. First, the main view was equipped with *telepointers* so that people could see each others' mouse cursor when they were looking at the same objects. Second, we provided *teleporting* as a fast way of looking at the other person's part of the workspace; pressing the right mouse button would immediately scroll the view to the other person's location, and remain there as long as the button was held down. This technique allows people to "glance" at another's work area without much effort. In addition to these basic features, systems could incorporate multi-user scrollbars, a wysiwid view, a miniature view, or a radar view.

The *multi-user scrollbar* (Figure 1F) shows each person's relative location in the workspace. The outside control is a standard scrollbar that allows people to move their own view. Beside the real scrollbar are indicator bars showing the current size and position of each person's viewport, each in a unique colour.

The "*what you see is what I do*" (wysiwid) view provides full-size details of another person's interaction, but shows only a limited part of their view (Figure 1D). The widget shows only the immediate context around another person's cursor, since most actions in graphical applications involve the mouse. As a person moves their cursor on a remote machine, the background of the widget pans to keep the display centred around the pointer.

The *miniature view* shows an overview of the entire workspace, reduced in area by a factor of 64 (Figure 1B). Each artifact in the workspace is represented in the miniature by a solid rectangle: text in grey, headlines in black, and pictures in different colours. Some information about others' activities is available through this display, since movement of or changes to any artifact are immediately duplicated in the miniature.

The *radar view* (Figure 1C) also uses a miniature view, but presents additional information about others' locations. The radar display shows the extent of what each person can see (their viewport) as a shaded rectangle, and also shows finer-grained location with telepointers that represent each person's mouse cursor. Participants are identified by showing view outlines and telepointer in unique colours.

Task

Each group was given two layout tasks, each with different artifacts. They were instructed to lay out the newspaper pages using their own knowledge of what newspapers should look like. Groups were allowed a maximum of fifteen minutes for each task. The pair was allowed to organize the completion of the task in any way they wanted, as long as they worked together; they were also asked to ensure that the two pages had consistent layouts, a constraint designed to prevent pairs from working totally independently. Since none of the subjects were experts in page layout, the task instructions were designed to create an informal atmosphere. Groups could ask the experimenters for assistance at any time.

Procedure

When subjects arrived, they were introduced to the experimenters, informed of the purpose of the study, and asked to sign consent forms. They were then given a guided tour of the system, its functions, and any awareness widgets that were part of the configuration. Participants then moved to their separate workstations and were allowed to practice each of the system's functions. When they felt comfortable, they began the first task. After fifteen minutes, the pair was stopped and given a questionnaire that explored their experiences with the system. Participants were then introduced to the second system configuration and the second task, also followed by a questionnaire. Finally, we conducted a short interview to investigate events that we had observed during the tasks and to explore particular responses on the questionnaires. Subjects were then paid and thanked for participating.

Study conditions

Six system configurations were used in the study (Table 1). Although pairs were randomly assigned to conditions, we did not equalize the group sizes within each configuration. We were interested both in getting a broad sampling of configurations and combinations, and in focusing on particular settings that we felt beforehand would provide useful information. Therefore, some widgets were assigned more subjects than others. The 'basic' condition shown in Table 1 refers to the system that included only the main view with telepointers and the teleport function.

Group	System 1	System 2
1	Basic	Radar
2	Radar	Basic
3	Mini	wysiwid
4	Scrollbar	Basic
5	Radar	Mini
6	Mini	Scrollbar
7	Radar	Radar + wysiwid
8	Radar	Basic

Widget	# Pairs
Basic	4
Scrollbar	2
wysiwid	2
Miniature	3
Radar	6

Table 1. System configurations and total pairs for each widget

Data collection

The experimenters recorded observations about the interaction and use of the system during the sessions. Questionnaires were administered after each task; these used 5-point Likert scales and balanced positive and negative questions. A short interview was also conducted and audiotaped. The task sessions were audiotaped and later transcribed (the transcripts were made after the analysis reported on here). A video record of one screen was taken to help construct the written transcript.

RESULTS AND INTERPRETATION

The study considered whether the widgets showed the right information, whether they were easy to interpret and make use of, whether they intruded on individual work, and how subjects felt about their value and their effects on the collaboration and final product. The results that we gathered, and some limited interpretation, are organized using these issues and are described below. We concentrate primarily on the four add-on widgets, but also consider the teleport function and main-view telepointers in some cases. We first present several general points that characterize our experiences in this study, and report on how the widgets were used.

General observations

All of the groups completed satisfactory layouts, and appeared to enjoy using the system. We observed a variety of collaboration styles, from top-down divide-and-conquer approaches where the group assigned tasks and then went off and did them, to close cooperation where the two subjects often worked on the same thing at the same time. In all of the pairs, manipulation of the artifacts was shared about equally between partners; we did not observe "single scribe" collaboration that has been noted in studies of collaborative writing [1]. Regardless of the style, all the groups used the relaxed-WYSIWIS capabilities of the system to focus at various times on separate areas of the workspace.

Even though the layout task did not force tightly-coupled interaction, it was clear that partners in all conditions maintained and used a sense of workspace awareness. We saw people gathering awareness information by asking their partner about where they were and what they were doing, or by watching them work. People also provided awareness information, primarily by verbally shadowing their own activities and intentions. We also observed utterances and actions that depended on awareness of location or activity, such as frequent deictic references to artifacts or areas of the workspace.

People did use the awareness widgets, and in general appeared to give them a fair trial. Subjects took enough interest that they gave us many comments and suggestions about the widgets, and their comments show that they thought carefully about how the widgets were, or could be, used. In some cases, people even became quite passionate about the widgets, especially when the second system did not provide a widget that they had liked using in the first!

The radar widget was found to be the most useful display of any that we tested. Both the miniature and radar widgets were considered to be much more useful than the multi-user scrollbars or the wysiwid view. Subjects gave more positive responses to the mini and the radar in both the questionnaires and in the interviews, and expressed their preference for them over the others when they had the opportunity to compare. Subjects reported that the miniature and radar were useful both because they assisted with a person's individual layout tasks, and because of the presentation of information about others in the workspace.

Amount and type of use

We first considered whether the widgets were used, and then looked at what they were used for. We regularly observed subjects making use of the widgets, and all of the subjects said that they made at least occasional use of the awareness widgets, if only to try them out. However, the amount of use varied across subjects, and some widgets were used more than others. Four of ten radar users commented that they didn't look at the radar much, and two users of the wysiwid display said they "never really used it." In some cases, people made little use of the displays because they forgot about them: one said "to be honest, I kind of forgot about the radar," and a user of the mini said "most of the time, I did not notice it." The teleport function was also regularly forgotten.

Part of this problem may have been caused by placement and visibility. The teleport function had no visual reminder of its presence, and the other widgets were small and were placed in the corner of the screen. One user said "the location of the radar was not very good. [It] forced you to look quite a distance."

When the widgets were used, it was generally in ways we expected. Primarily, they were used to gather information about the other person in the workspace. However, the mini and radar displays were also used to aid individual work: for example, people often looked at the mini or the radar to manipulate text columns that were too long to be seen completely in the main view. The wysiwid view was used in more limited ways: one person said that when they knew "that my partner was on a picture… I would use the wysiwid to determine *which* picture."

Information content

Our first goal was to explore whether the widgets collected and presented the right information. We discussed several issues of information content with the subjects: amount and type of information presented, whether the displays adequately showed location and activity in the workspace, and whether there was any information missing from the widgets. We found that subjects responded positively to widgets that provided more information; however, they also thought that additional types of information should be available.

Responses to questions about amount and type of information were consistently more positive for the systems that provided additional awareness information. That is, the basic configuration was ranked lowest, the scrollbars were given a neutral response, and then increasingly positive responses were given to the wysiwid, the mini, and the radar. The subjects' responses suggest that they saw a lack of information in the basic configuration, and that the information presented in the widgets provided at least some of what was missing. This pattern was repeated for questions that asked specifically about whether the widgets adequately showed the location and activities of the other person. Expected variations were recorded, however, for widgets that did not provide particular information (for example, the scrollbars do not show activity, and the mini does not show view location). The high marks given to the radar view suggest that it provides approximately the kind and amount of information needed to perform the layout task.

However, subjects also mentioned additional kinds of information that they thought should be presented by the systems. In various configurations, subjects felt that information was missing about their partner's location, activities, intentions, the overall look of the workspace, and about the details of the other person's view. In some cases, this information would have been available through other widgets. For example, the basic configuration was often found to be lacking in location information (e.g. it was "difficult to tell at times where my partner was working") and in information about the overall workspace (e.g. "[it was] hard to tell [the] layout of the page;" "[it] would be nicer with a full view of the 2-page area."). This information could have been provided by the radar view, and two subjects using the basic system came up with the

idea of miniature and radar views as ways to address these deficiencies.

Some of the requests, however, concerned information about type of activity and intentions that was not available in any of the widgets. Four subjects felt that information about "the type of operation being performed by my partner" was missing. Two people specifically mentioned that they could not tell when their partner had grabbed an object to move it, information that would have warned someone against trying to move the same object (which happened occasionally). Also, people wanted to know when their partner was going to cut a column of text, an activity that occasionally hung the system if there was too much other activity going on. Subjects also wanted information about their partner's intentions. A few comments concerned "what my partner was going to do next" and their "future plan." Subjects agreed that this would be difficult for the system to determine, but some people did make concrete suggestions. For example, one subject wished that he had been able to mark the objects that he was planning to use in the near future, so that his partner would know not to take them.

Ease of interpreting the information

Once we had considered the information content of the widgets, we wanted to find out how easily the subjects could interpret and make use of that information. We examined several issues: how easy it was to get location and activity information from the displays, whether subjects could determine identity in the widgets, and whether the displays were confusing. We found that for each of these, the radar and mini displays were easier to interpret than the scrollbars or the wysiwid view.

Information about the other person's location was generally considered easier to interpret in the radar and mini views than in the other widgets. There are several reasons why the wysiwid view and the scrollbars caused problems. The wysiwid's animation was not smooth, which may have made it difficult to understand. One user said "identifying where your partner is hard because my partner's movement in [the wysiwid] is very jerky." Subjects' responses to questions about the multi-user scrollbar indicated that although the scrollbars presented the relevant information, it was difficult to actually determine where the other person was. This may have more to do with integration of information than with any technical difficulties. The scrollbar showed each user's view location, but split the extents of the rectangular view into horizontal and vertical dimensions. To determine where someone was, a user had to mentally integrate information from the two scrollbars. The radar widget, in contrast, showed view location in a form that was already integrated.

Unlike location, the subjects felt that in no cases was it particularly easy to interpret what their partner was doing. One contributing factor may be that the widgets did not display the type of activity that the other person was undertaking. However, this response may also indicate the inherent difficulty of determining high-level task goals from watching low-level activity, such as that presented in the widgets.

Subjects had little difficulty in determining who was who in the widgets. However, since there were only two people in the workspace, they often did not need the colour cues that the widgets use to distinguish participants. Several people looked for motion in the widgets to identify their partner, rather than remembering who was represented by each colour. For example, one subject determined his partner's location by "wait[ing] to see what was being done on the screen that was not connected to my actions."

There were also few reports of difficulty in switching between the context of the main view and that of the widgets, despite the fact that two widgets showed the workspace at a different scale, and two showed different parts of the workspace. A few problems were encountered with adjusting to the wysiwid view, and one subject suggested that it showed too little of the other user's view for him to figure out where it was. The general lack of difficulty, however, may be due to the use of representations that were similar to the main view, and therefore familiar to the subjects. The artifacts in the workspace could act as landmarks to help tie the two representations together, and could therefore help people determine (for example) someone's actual location based on the position of their telepointer in the radar.

Effects on individual work, collaboration, and product

We were also interested in how the widgets affected individual activities, interaction, and the final layout. We found that the subjects were not distracted by the mini or radar widgets, and that several people felt the radar to have a significant positive effect on their work. However, we did not see obvious differences in collaboration or final product, and some subjects had problems in acting on the information that they had gathered.

Subjects stated strongly that the mini and radar views did not distract them or slow down the completion of the task. One subject said "the mini display was not distracting in the least," and there is added evidence of the widgets' unobtrusiveness in the fact that several subjects forgot about the displays altogether! Subjects were more ambivalent about the scrollbars and the wysiwid view. Two subjects found that the scrollbars distracted them from their tasks, and another considered the wysiwid "almost distracting" because it "shows too small an area to gain any real benefit from it."

Five subjects also expressed the belief that the radar and mini views did have a significant positive effect on the final outcome. For example, subjects said "we really needed the

radar to help in the overall appearance," "the final result would have been much worse without the radar display," and "the radar screen made the task possible... I think without it would have been difficult to complete the task." There was, therefore, at least a perceived effect on the outcome.

Despite subjects' enthusiasm for the radar and mini views, we did not see that particular configurations made obvious differences in the collaboration or final product. All pairs were able to produce reasonable layouts in fifteen minutes, and subjects responded positively to all systems when asked if they allowed efficient and effective collaboration. Groups with unaugmented systems did not seem to have more difficulty in coordinating or carrying out their tasks than groups with the mini or radar widgets. This may be due simply to people's ability to adapt their behaviour in information-poor situations: people are remarkably capable of finding information that they need through alternate means, and they can often find ways to complete their tasks without the missing information. However, this particular result is preliminary, and differences in collaboration and performance may surface as we undertake more detailed analyses of verbal protocols.

A few subjects noted problems in the relation between getting information from the displays and putting it to use. One subject said "it was easy to get info about my partner but the 'bridge' between getting info *about* and then proceeding to go over to *where* my partner was, is slow." In the situation he was describing, he had seen his partner do something in the wysiwid view that he wanted to help with, but it took him a long time to scroll his main view to where she was. The problem arose because the widgets in the study were designed only to provide information, not to help people act on that information. Several other participants added to this by suggesting that the widgets should be manipulable: people asked to be able to move artifacts in the mini view, and to be able to change the location of their main view by dragging their view rectangle in the radar.

Perceived value

Despite the lack of obvious effects, it was clear that the participants in the study found some of the widgets valuable. We asked subjects whether the widgets they used were valuable additions to the groupware system, and whether the widgets were worth the screen space that they took away from the main view. Subjects felt strongly that the radar was a valuable addition to the system, and that the radar and mini displays were worth their screen space. When comparing the two system configurations that they had used, subjects almost always preferred configurations incorporating these widgets. We often received joking complaints when the second system "took away" the radar view. In addition, several people said that the radar view was important to the completion of the layout task, and there were numerous positive comments on the radar and mini displays. For example, different subjects said that they were "very helpful," "extremely helpful," "a must," and "a very useful tool in groupware applications." Responses were mixed for the other widgets, and it appeared that participants did not value the scrollbars and the wysiwid view as much. However, several people made suggestions about improving these widgets, or speculated on situations where they would be more useful, indicating that they believed even these devices had some promise.

In the interviews, we also explored the reasons why the mini and radar displays were useful. Responses indicated that the value of the widgets was about evenly divided between their ability to assist individual work, and the workspace awareness information that they provided. That people found the awareness information valuable, above and beyond the overview, was also shown in the comments of three subjects who suggested that the mini display could be improved by showing their partner's telepointer.

As a final point, a comment made by one subject has much to say about the promise of displays like the radar view for making virtual workspaces more natural. When asked about the widget's value, he said "it really felt like you were working on the same big table."

DISCUSSION

Here we consider possible explanations for our results, discuss the implications for groupware designers, compare our findings to previous work, and comment on possible limits to the generality of our conclusions. A recurring theme in the following paragraphs, however, is the relationship between the information shown in the widgets and the spatial constraints of the task.

Explanations of results

The main results of the study are that the radar and mini views were fairly successful in helping people maintain workspace awareness in a page-layout task, and that the multi-user scrollbar and the wysiwid view were not. There are several possible reasons for this outcome, and these are organized below in terms of information and interpretation.

Information content. There were two ways that the mini and radar provided better information than did the wysiwid view and scrollbar. First, the mini and radar views provided more types of information that were important to the layout task. They showed both the overall workspace and the other person's interactions, whereas the wysiwid and scrollbar showed only the details of activity and location respectively. The scrollbar in particular shows only view extents, which may simply have not have been enough information. Second, the mini and radar presented information at a more appropriate level: showing location and actions in terms of the overall workspace appeared to fit the way that groups visualized the task better than the lower-level presentation of the wysiwid view.

The success of the miniature view (which did not show view rectangles or telepointers) may seem to imply that the workspace overview was the primary reason for the success of both the mini and the radar. However, two things suggest that awareness information was also important. First, subjects made use of the considerable awareness information was presented by the miniature: by watching artifacts move, subjects could determine both the location and activity of the other person. Second, there were distinct preferences for the radar when it was compared to the mini. The major advantage of the radar view appeared to be its telepointers, which showed location even when nothing was being moved, and allowed gesturing.

Ease of interpretation. The radar view was found to be easier to interpret than the multi-user scrollbar, and comparing these two widgets illustrates the benefits of familiar representations for presenting awareness information. Subjects were able to use their existing knowledge of the workspace when interpreting awareness information in the radar view. The obvious similarity between the main workspace and the radar's scaled representation provides a familiar base on which to present other spatial information (such as view extents and cursor locations). The differences in scale and in representation detail (i.e. objects were represented by boxes) did not appear to be a problem, suggesting that shape, relative size, and area provide sufficient cues for mapping between the radar and main views.

In contrast, the multi-user scrollbar did not present awareness information on a familiar base: it shows view location as a range (the indicator bar) on an abstract scale (the length of the scrollbar), which has no spatial connection to artifacts or distances in the actual workspace. This presentation limits the information that can be drawn from the scrollbar to an understanding of whether two people share some part of their view, and unless the two indicators are perfectly aligned, it is difficult to determine exactly which artifacts can be seen by both people.

The overview was particularly useful in this task, where seeing the overall layout was important. However, we suspect that overviews will prove valuable in other types of shared-workspace tasks as well, since workspace artifacts can provide landmarks for interpreting information about location and activity. In addition, because other participants can alter the workspace without one's knowledge, overviews can help people keep track of things that they may later need to find.

As previously mentioned, a related problem with the multi-user scrollbar is that it splits the presentation of a view rectangle into two one-dimensional displays, forcing the user to reintegrate the information. The two-dimensional presentation of the radar view, again, allows people to use their existing models of the workspace.

Lessons for groupware designers

This study has shown that support for workspace awareness can be a valuable and appreciated addition to a groupware system. Groupware designers should consider requirements for workspace awareness when they design shared-workspace systems, and can use add-on widgets to help people maintain awareness.

In particular, workspace overviews as used in both the radar and mini views can assist people in tasks where spatial manipulation and spatial relationships are important. If tasks require only general knowledge of others' locations and activities, the wysiwid display is likely to be ineffective, since it shows only the details of others' actions. Also, since integration of the spatial information is critical, it is likely that the multi-user scrollbar is not a good candidate for use in two-dimensional workspaces.

In addition to recommendations about particular widgets, our experiences suggest that future studies of groupware should include criteria for assessing awareness. Standard usability evaluation methods generally do not assess support for the group dynamics of a shared workspace system. We propose that criteria such as information content and ease of interpretation be used more generally to consider how groupware provides information about other participants.

Comparison to previous work

A number of this study's conclusions can be contrasted with previous CSCW research. In particular, we consider previous usability studies of awareness widgets, the issue of distraction, and the principles of passive collection and workspace presentation of awareness information.

Baecker and colleagues [1] report on the only other usability study that we know of to include awareness widgets. They evaluated a synchronous text editor that included a multi-user scrollbar and a text overview similar to the radar view. They do not provide detailed results, but say that "most subjects were able to use SASSE's... collaborator awareness mechanisms successfully" (p. 404). This stands in contrast to our subjects' difficulties with the multi-user scrollbar. However, this widget is likely more suited to a text editor than it was to our system, since text usually has only one dimension, and since relative positions may more closely represent the way people think about location in text documents.

Ellis and colleagues [8] considered the issue of distraction in groupware interfaces, and state that these interfaces should present information but not be overly distracting. Our study, in which the radar view was not considered distracting, leads us to some speculation on this issue. It is commonly thought that distraction is caused by perceptual information that draws our attention, but distraction may have as much to do with interpretation difficulty. If people already have a good sense of workspace awareness, then

changes in a display like the radar are expected and will generally not distract. Ellis et al. state that groupware participants "are generally not as aware of others' contexts and can less easily interpret sudden display changes resulting from others' actions" (p. 49). To this we add that if people can stay aware of others' contexts, then display changes will be easier to interpret and thus less distracting.

Dourish and Bellotti [7] give two principles for supporting awareness in groupware workspaces. One is that awareness information should be "presented in the same shared work space as the object of collaboration" (p. 107). The widgets in this study show, however, that awareness information can be successfully presented outside the main view. Other systems do so as well: for example, Cognoter [22] indicated activity in the icon of a closed window, and several systems have presented sound cues separate from any workspace (e.g. [4,10]). In our view, a more general principle should be used: that awareness information must be easily interpretable regardless of where it is presented. We agree, however, that one way of simplifying interpretation is to build on people's existing knowledge of the workspace.

The other principle stated by Dourish and Bellotti is that awareness information should be passively collected and distributed by the system, rather than explicitly generated by the participants. Again, we agree with the principle in general; however, two episodes in our study suggest that there are situations where it does not always hold. First, we regularly observed people verbally shadowing their own behaviour (e.g. "now I'm moving the picture up to the top..."), an activity that we and others have seen in other group situations (e.g. [15]). People take on the task of shadowing quite naturally, without any prompting. It may be an indication that people are interested in helping their collaborators maintain awareness. In the second episode, a subject suggested that he should be able to mark objects in the workspace in order to indicate to his partner that he intended to use them in the near future. This is an example of information that cannot be passively gathered by the groupware system, and suggests that some awareness information can only be generated explicitly by the participants.

Critical reflection

There are a variety of issues that may limit the generality of our conclusions. The two most important concern group size and task type.

The first issue considers how the widgets will scale to larger groups. Many of our results may change somewhat if more than two people are in the workspace. It is possible that the increasing amount of awareness information in the radar would begin to clutter the widget. The one subject who thought that two view rectangles cluttered the display might be joined by several others when there are four rectangles. In the mini view, subjects will have far more difficulty in determining who is moving an artifact, since they can no longer adopt the "if it's not me, then it's my partner" identification strategy. Although the radar shows different coloured telepointers that can assist in identifying people, this technique forces users to map colours to people. The wysiwid display is not itself affected by additional people, since it shows only one person; however, users would have to decide whom they wish to track using the display, or sacrifice more screen space to have multiple wysiwid views. Although we believe that none of these concerns would render the widgets useless in a larger group, it is obvious that they must be examined in more detail.

The second issue concerns how our results would change in tasks other than page layout. The task we used involved a medium-sized workspace where participants could see only about a third of the total space. The task was completed through spatial manipulation of the artifacts; in addition, spatial relationships between artifacts and the overall appearance of the workspace were important criteria in completing the task successfully. Radar and mini widgets should thus be useful whenever an overview helps a person maintain a model of spatial changes in the workspace, as well as awareness of how the changes are being made. We believe our results will hold for other tasks that share these characteristics, and that they apply in limited ways to other artifact-based tasks as well. In other tasks, other kinds of constraints may require that overviews be organized differently. For example, a text document may be better served by an outline overview rather than a miniature representation of every line.

Future work

We plan to build on this work in several ways. First, we will incorporate some of the subjects' suggestions into new versions of the widgets. In particular, we will improve the bridge between perception and action that one subject alluded to. Second, we will address the concerns raised in the previous section, and test the radar view with larger groups and on other kinds of tasks. Third, we hope to undertake further studies that use stronger measures of the widgets' effectiveness, such as quality of product or time to completion. Finally, we also plan to test new designs, such as a fisheye workspace that uses multiple focal points to convey workspace awareness information [12].

CONCLUSIONS

This paper has described a usability study of awareness widgets added on to a shared-workspace groupware application. Two displays based on miniature overviews of the workspace were particularly useful, both for individual tasks and for maintaining workspace awareness. We expect that the issues raised here will motivate groupware designers to continue exploring awareness, with the goal of building shared workspaces that are as natural to use as their physical counterparts.

Acknowledgments

We are grateful to Intel Corporation, NSERC, and the Province of Alberta for financial assistance.

REFERENCES

1. Baecker, R., Nastos, D., Posner, I., and Mawby, K. The User-Centred Iterative Design of Collaborative Writing Software. In *Proc. INTERCHI'93* (Amsterdam, 1993), pp. 399-405.
2. Beaudouin-Lafon, M., and Karsenty, A. Transparency and Awareness in a Real-Time Groupware System. In *Proc. UIST'92* (Monterey, CA, 1992), pp. 171-180.
3. Borning, A. and Travers, M. Two approaches to casual interaction over computer and video networks. In *Proc. CHI'91* (New Orleans, 1991).
4. Cohen, J. Monitoring Background Activities. In Kramer, G. (ed.) *Auditory Display*. Addison-Wesley, 1994, pp. 439-531.
5. Conklin, J. gIBIS: A hypertext tool for exploratory policy discussion. In *Proc CSCW'88* (Portland, OR, 1988), pp. 140-152.
6. Dourish, P., and Bly, S. Portholes: Supporting Awareness in a Distributed Work Group. In *Proc. CHI '92* (Monterey, CA, 1992) pp. 541-547.
7. Dourish, P. and Bellotti, V. Awareness and Coordination in Shared Workspaces. In *Proc. CSCW'92* (Toronto, 1992), pp. 107-114.
8. Ellis, C., Gibbs, S. and Rein, G. Groupware: Some Issues and Experiences. *CACM 34*, 1 (1991), 38-58.
9. Englebart, D. and English, W. A Research Center for Augmenting Human Intellect. In Greif, I. (ed). *Computer-Supported Cooperative Work*. Morgan Kaufmann, San Mateo, 1988.
10. Gaver, W., Smith, R., and O'Shea, T. Effective Sounds in Complex Systems: The ARKola Simulation. In *Proc. CHI'91* (New Orleans, LA, 1991), pp. 85-90.
11. Gaver, W. Sound Support for Collaboration. In *Proc. ECSCW'91* (Amsterdam, 1991), pp. 293-308.
12. Greenberg, S. and Gutwin, C. Sharing Fisheye Views in Relaxed-WYSIWIS Groupware Applications. To appear in *Proc. Graphics Interface '96* (Toronto, 1996).
13. Gutwin, C., and Greenberg, S. Workspace Awareness for Groupware. To appear in *Proc. CHI'96* (Vancouver, 1996).
14. Gutwin, C., Greenberg, S., and Roseman, M. Workspace Awareness Support With Radar Views. To appear in *Proc CHI'96* (Vancouver, 1996).
15. Heath, C., and Luff, P. Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. In *Proc. ECSCW '91* (Amsterdam, 1991), pp. 65-80.
16. Roseman, M., and Greenberg, S. Designing Real-Time Groupware with GroupKit, a Groupware Toolkit. To appear in *Trans. on CHI* (1996).
17. Short, J., Williams, E., and Christie, B. Visual Communication and Social Interaction. In *The Social Psychology of Telecommunications*. John Wiley & Sons, 1976, pp. 43-60.
18. Smith, R., O'Shea, T., O'Malley, C., Scanlon, E., and Taylor, J. Preliminary Experiments with a Distributed, Multi-Media, Problem Solving Environment. In *Proc. ECSCW'89* (1989), pp. 19-34.
19. Stefik, M., Bobrow, D., Foster, G., Lanning, S. and Tatar, D. WYSIWIS Revised: Early Experiences with Multiuser Interfaces. *TOIS 5*, 2 (1987), 147-167.
20. Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S. and Suchman, L. Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. *CACM 30*, 1 (1987), 32-47.
21. Tang, J. C. Findings from observational studies of collaborative work. *IJMMS* 34, 2 (1991), 143-160.
22. Tatar, D., Foster, G. and Bobrow, D. Design for conversation: Lessons from Cognoter. *IJMMS* 34, 2 (1991), pp. 185-209.