Awareness Through Fisheye Views in Relaxed-WYSIWIS Groupware

Saul Greenberg and Carl Gutwin
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4

phone: +1 403 220 6087
email: saul *or* gutwin@cpsc.ucalgary.ca

Andy Cockburn
Department of Computer Science
University of Canterbury
Christchurch, New Zealand

phone: +64 3 364 2774
email: andy@cosc.canterbury.ac.nz

Abstract

Desktop conferencing systems are now shifting from strict view-sharing towards relaxed "what-you-see-is-what-I-see" interfaces, where distributed participants in a real time session can view different parts of a shared visual workspace. As with strict view-sharing, people using relaxed-WYSIWIS require a sense of *workspace awareness*—the up-to-the-minute knowledge about another person's interactions with the shared workspace. The problem is deciding how to provide a user with an appropriate level of awareness of what other participants are doing when they are working in different areas of the workspace. In this paper, we summarize requirements for workspace awareness, identify limitations of existing groupware solutions, and propose as a replacement *fisheye views* that show both global context and local detail within a single window. Within groupware, these displays provide: a) peripheral awareness of other participants by showing their position and actions in the global context; b) detailed awareness of interactions by assigning multiple focal points for each participant, and by magnifying the areas where they are working. Two groupware prototypes illustrate these concepts: a fisheye graph browser, and a fisheye text viewer.

Keywords: fisheye views, awareness, groupware, information visualization.

1. Introduction

Real-time distributed groupware helps people who are geographically separate to work together at the same time (Baecker 1993). These networked computer applications provide a shared virtual workspace where people can see and manipulate work artifacts, much as face-to-face work often occurs over a shared physical workspace like a whiteboard, text document, or control panel. Virtual workspaces support various group activities, such as shared drawing (e.g. Greenberg, Hayne and Rada 1995), shared text editing (e.g., Baecker, Glass, Mitchell and Posner 1993), idea generation and organization (e.g., Tatar, Foster and Bobrow 1991; Valacich, Dennis and Nunamaker Jr 1991), or multi-user games. In addition to the workspace, a groupware system will likely incorporate facilities for communication over audio and video links.

The problem is that groupware workspaces cannot yet match the diversity and richness of interaction that their physical counterparts afford. In particular, virtual workspaces make it more difficult to maintain a sense of awareness about who else is in the workspace, where they are operating, and what they are doing. In a physical workspace, people use peripheral vision, auditory cues, and quick glances to keep track of what goes on around them. In a groupware system, the visual field is greatly reduced, and many of our normal mechanisms for gathering information (such as glancing) are ineffective since the required information may be absent from the display.

In addition, the way that a groupware system supports view sharing can further impair people's abilities to stay aware. Recent groupware systems have relaxed the strict "what you see is what I see" (WYSIWIS) model (Stefik, Bobrow, Foster, et al. 1987), where all participants see exactly the same view of the workspace at all times. The relaxations give people control over their own viewport into a workspace, and thus allow them to work in a more natural style, shifting their focus back and forth between individual and group work. Relaxed-WYSIWIS view-sharing makes groupware more flexible and better matches the way people actually work, especially in large workspaces that contain many artifacts. However, relaxed-WYSIWIS can also contribute to loss of awareness since, when views differ, people can lose track of where others are and what they are doing in the workspace.

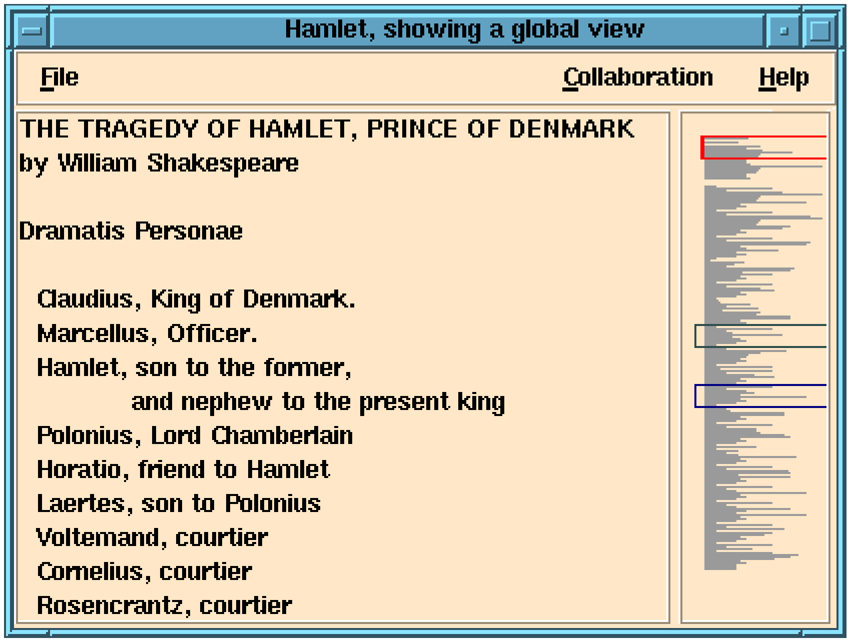

Figure 1. GroupKit's radar view widget (Roseman and Greenberg 1996). The miniature on the right shows the entire document. Viewports of participants are overlaid as coloured boxes, with the current participant's viewport at the top. The user can scroll their viewport by grabbing and moving their box.

Several approaches for providing awareness cues in relaxed-WYSIWIS systems have been proposed (e.g. Gutwin and Greenberg 1995; Gutwin, Greenberg and Roseman 1996). One popular method supplies users with two separate windows, one containing a detailed view and the other a "radar" overview (e.g. Smith, O'Shea, O'Malley et. al. 1989; Baecker, Glass, Mitchell and Posner 1994). Figure 1 shows a sample radar view display, supplied as a widget within GroupKit, our groupware toolkit (Roseman and Greenberg 1996). The *detail window* on the left is a full-sized view of part of the workspace—in this case a text document—and is where a person does their work. The *radar window* on the right is an overview window presenting a miniature of the workspace, typically overlaid with boxes that represent each participant's viewport. The radar view may even show others' actions and telepointer movements as they occur, although at very low resolution.

Although radar views do provide support for maintaining workspace awareness (Gutwin, Roseman and Greenberg 1996), they have two major limitations.

1 Radar views introduce a physical and contextual gap between local and global contexts. In order to gather awareness information from the radar view, people must move their attention to a different part of the screen and make an abrupt context shift to the scale and extents of the overview presentation. This gap does not exist in face-to-face interaction over physical workspaces, where perceptual abilities imply a gradual loss of detail for more distant objects. A related problem with radar views is that people have to mentally integrate the two displays, matching their detailed view with the boxes in the overview, and determining how other people's viewports relate to theirs.

2 Radar views do not support lightweight mechanisms for maintaining awareness of detailed activity. The low resolution of the radar overview obscures the details of another person's actions. To determine exactly what someone else is doing, one has to align viewports, usually by scrolling their own detailed view. This involves far more effort than a simple glance over to another part of a physical workspace.

We propose using *fisheye views* to overcome these two limitations. First, fisheye representations of a workspace provide a seamless and smooth transition from local to global contexts, providing a more natural scene in which to provide awareness information. Second, fisheye views allow the use of multiple focal points, one for each participant, which can provide improved awareness of the details of others' actions in the workspace.

In this paper, we first describe what is meant by *workspace awareness* and explain why it must be supported in relaxed-WYSIWIS groupware systems. Section 3 then presents a brief background on fisheye views and discusses how they can support awareness. Two prototype fisheye displays are then presented as case studies. The paper closes by evaluating the strengths and weaknesses of these fisheye views as techniques for maintaining workspace awareness.

2. Workspace Awareness

Awareness about who is present in the workspace, where they are, and what they are doing, is all part of what we call *workspace awareness*—the up-to-the-minute knowledge about another person's interactions with the shared workspace. In face to face shared activity, workspace awareness is a natural, constant, and even unconscious part of people's interaction. In *tightly coupled* collaboration, where participants interact closely, awareness is maintained through speech, through observation of others' actions in the workspace (Segal 1994), through gestural communication (Tang 1991), through deictic references (Tatar, Foster and Bobrow 1991), and through observation of the direction of another's gaze (Ishii and Kobayashi 1992). In *loosely coupled* collaborations, where the group pursues separate but still co-ordinated goals, awareness is maintained by peripheral vision, by quick glances at others' areas, and by brief utterances that inform others of activities and intentions (Gutwin and Greenberg 1995; Heath and Luff 1991). Groups often combine these two kinds of coupling and work in a *mixed-focus* fashion, where participants move back and forth between shared and individual work. In mixed-focus collaboration, workspace awareness is particularly useful in helping people to manage these transitions and recognize opportunities for closer collaboration (e.g. Dourish and Belloti 1992).

Workspace awareness is one of several overlapping types of awareness that people maintain when they work in a group, as shown in Figure 2.

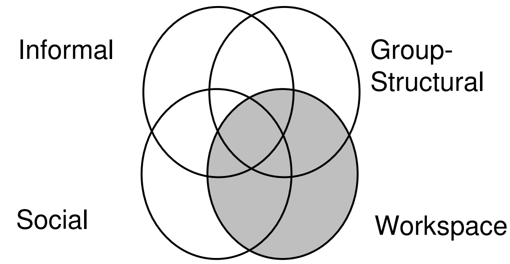

Figure 2. Four types of awareness in collaborative work

- *Informal awareness* of a work community is the general sense of who's around and what they are up to—the kinds of things that people know when they work together in the same office. Informal awareness facilitates casual interaction, and is often supported in groupware through media spaces (e.g., Bly 1993).
- *Social awareness* is the information that a person maintains about others in a social or conversational context: whether another person is paying attention, their emotional state, or their level of interest. It is maintained through back-channel feedback, and through non-verbal cues like eye contact, facial expression, and body language. It is typically supported in groupware through desktop video conferencing (e.g., Ishii and Kobayashi 1992).
- *Group-structural awareness* involves knowledge about such things as people's roles and responsibilities, their positions on an issue, their status, and group processes. Here, groupware supports the group by making the group's structures, processes, and roles explicit (e.g. Valacich, Dennis and Nunamaker Jr 1991; Leland, Fish and Kraut 1988).

The fourth type is workspace awareness, different from the others because of the workspace's integral role in the collaboration. While it is less easy to define exactly what knowledge people require, the first column in Table 1 (in the Discussion section) summarises a few of the more essential elements comprising awareness, phrased as questions (the framework is fully described in Gutwin, Stark and Greenberg 1995). These awareness factors include information on the following important items: the *identity* of those in the workspace, their *location*, their *activity*, and the immediacy of *changes* with which others' activities are communicated. The elements in this table provide guidelines for the development of the prototypes described in later sections.

Workspace awareness has also been recognized in groupware research (although under different names), and our work builds on these efforts (Dourish and Bellotti 1992; Bly, 1993; Beaudouin-Lafon and Karsenty 1992; Baecker, Glass, et al. 1994). In addition, our investigations of support for workspace awareness is based on two key differences between face-to-face and groupware situations.

- The perceivable environment is drastically reduced in groupware. In face-to-face interaction, people can generally see the entire physical workspace and all the people in it; in groupware, they have only a small window into the virtual space. The reduction of the visual field is one reason that relaxations to the WYSIWIS model have been proposed, increasing individual control but decreasing support for workspace awareness.

- The differences between being immersed in an environment and having only a window onto it imply that many of the natural mechanisms that we use to gather awareness information in a face-to-face setting (such as peripheral vision or glancing) will be inappropriate when sitting in front of a computer screen. Computational analogues to these mechanisms (such as scrolling the viewport) are slow and clumsy in comparison.

Within this strange new situation, the groupware designer must try and recreate the conditions and cues that allow people to keep up a sense of workspace awareness. Whereas face-to-face interaction has inherent mechanisms and affordances for maintaining awareness, the groupware designer is faced with a blank slate—any support for building or maintaining workspace awareness must be explicitly chosen and built into the groupware system. Our research involves the design of groupware that supports some of people’s awareness needs (Gutwin and Greenberg 1995; Gutwin, Stark and Greenberg 1995; Gutwin, Greenberg and Roseman 1996), and the construction of *awareness widgets* for our groupware toolkit, GroupKit (Roseman and Greenberg 1996). Recently, we have designed and built two displays that are based on the idea of fisheye views. These widgets begin to address the two key issues above, and show promise for extending current approaches to the support of workspace awareness. The following section describes fisheye views, and then introduces these two widgets and the ways that they support awareness.

3. Applying Fisheye Views to Groupware

A brief background to fisheye views. Fisheye views are computer visualization techniques that provide both local detail and global context in a single display. They take their name from the photographer's fisheye lens, a hemispherical lens that distorts a scene to provide an extremely wide angle of view. In a computational fisheye, the user chooses a point of focus where they wish to see local detail: this area is visually emphasized, and the remainder of the data is made less visually important.

Fisheye views have been used to visualize data in many domains. Furnas (1986) created systems for viewing and filtering structured program code, biological taxonomies, and calendars. Egan, Remde, Landauer et. al. (1989) used a type of fisheye view in Superbook, a text-based electronic book, to provide the now familiar notion of an expandable table of contents. Sarkar and Brown (1992) implemented graphical fisheye views for networks of nodes such as cities on a map.

Computational fisheye views generally use combinations of three different presentation techniques (Schaffer 1995), shown by way of example in the organization chart in Figure 3. First, the standard representation (Figure 3a) can be graphically distorted to enlarge the area around the focus point (the ‘president’ node) and reduce the size of more distant objects (Figure 3b). The size of an object is determined by its distance from the focus point, but size can also be altered depending on the node's importance or interest value (Furnas 1986; Sarkar and Brown 1992). Second, fisheye views can provide more space for detail in the focus area by replacing distant objects with simpler representations that use less room. For example, the view in Figure 3c clusters objects outside the focus by replacing the full text descriptions with labels. Third, fisheyes can filter the data space, removing low-interest items from the view altogether. In Figure 3d, for example, the lower elements of the original chart are not displayed, leaving more room for the other elements. Sarkar and Brown’s graph browser (1992) is another example of importance-based filtering.

The above methods alter the representation of a data space to emphasize a single focus, but multiple focus points can also be supported with fisheye views. Sarkar, Snibbe, Tversky and Reiss (1993) built displays based on the metaphor of a rubber sheet, where several different focal points could be “pushed forward” for emphasis. In addition, this system gave the user direct control over the amount of screen space used for objects in the areas of interest. Schaffer, Zuo, Greenberg et al. (1996) also provided for multiple focal points in hierarchically-clustered networks.

The displays that we have been building use graphical distortion and the idea of multiple focal points as a basis for supporting workspace awareness in groupware systems. In the following sections we present two widgets. The first illustrates how location information can be provided in a fisheye network browser. The second is a more sophisticated fisheye for viewing text files, and incorporates support for location awareness, multiple focal points to show details of others’ activity, and customizable lenses to change the amount of space given to the local or the remote focal points.

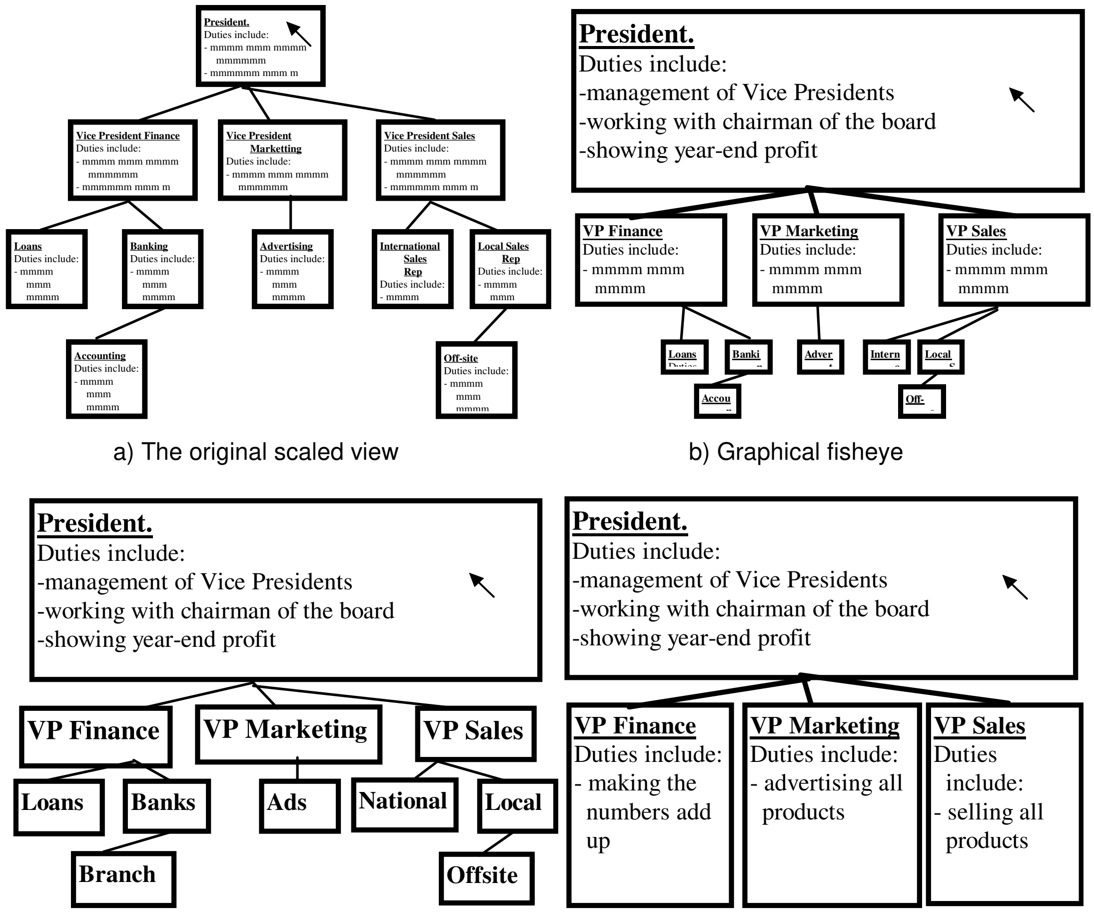

a) The original scaled view

b) Graphical fisheye

c) Outline fisheye

d) Filtered fisheye

Figure 3. Applying various fisheye view techniques to an organizational chart. The focal point is on ‘President’

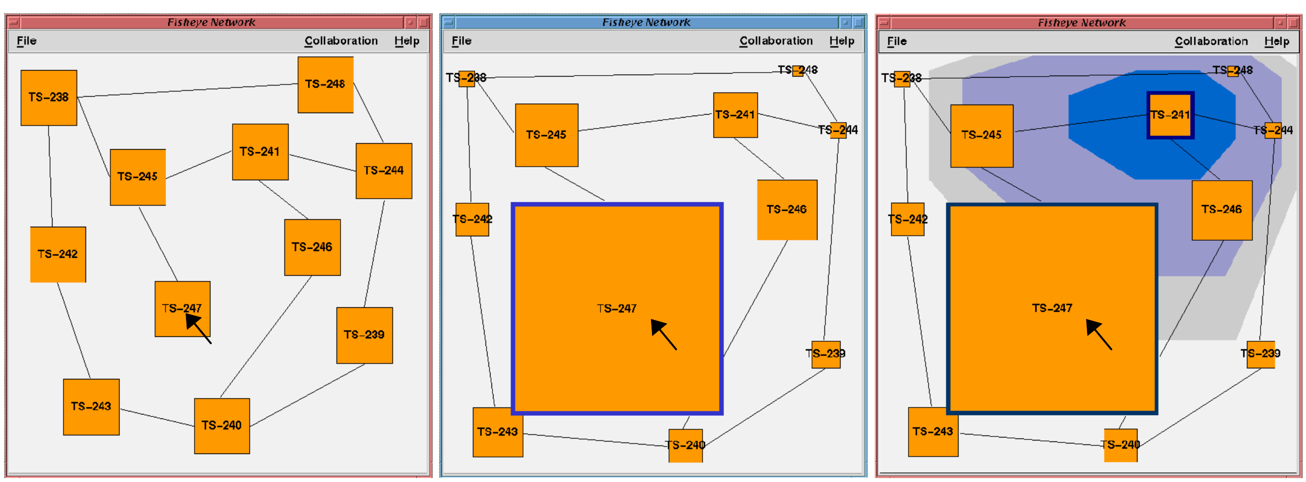

a) undistorted network

b) single user fisheye view

c) multi-user fisheye view

Figure 4. A simulated network showing its original state, its fisheye view, and how location of others can be displayed in a multi-user view through outlining and halos

Location awareness in a graphical fisheye view. Our first prototype exploits a fisheye view's ability to represent an entire workspace, thus allowing information about other participants' locations to be displayed regardless of where others are working. The prototype operates on a crude two-dimensional network of nodes, such as the (undistorted) example shown in Figure 4a. In order to show the entire data space, the nodes in the undistorted view must be reduced in size to the point where it is difficult to see detail.[1] This is remedied by applying a fisheye lens over a focal point (Figure 4b), implemented with the algorithm described by Sarkar and Brown (1992). Here, the nodes in the network around the focal point (the large, outlined box) are seen in detail, while still retaining global context.

In a relaxed-WYSIWIS groupware version of the browser, each person has control over their own focal point, and can therefore see details of the parts of the network that are required for their tasks. Consider the example in Figure 4c. The local person is still focused on the node just southwest of center, while their partner is viewing a node to the north-west. Each person's display will thus be "fisheyed" differently. However, the viewer still shows information about location and focus, helping to keep the group aware of one another's whereabouts and of what each other are looking at. First, the node at each remote person's focus point is outlined with a "personal" colour—this is shown in Figure 4c by the outlined box on the top right. Second, a diminishing "halo" extends outwards from the remote person's focal point that corresponds roughly to that person's area of focus on their own displays.

In undistorted space, halos are normal eight-sided polygons. To represent the halos in each participant's fisheye view, however, these are mapped appropriately to the particular spatial distortion that has already been applied to the data. The result are irregular polygons seen in Figure 4c.

Although this prototype can convey information about other participants' locations and how their focus changes over time, it does not show any details of their actions. To better support awareness of others' activity, we have constructed another prototype that uses multiple focal points to show details of activity as well as location information.

Location and detail awareness in the fisheye text viewer. Our second prototype is the fisheye text system which is both a groupware text viewer and rudimentary groupware editor. It too indicates people's locations within a workspace, but also illustrates how details of other people's activity can be presented via multiple focal points. To demonstrate the fisheye text viewer, we first show how it works with a single-user, and then present it as a multi-user system.

The viewer uses a fisheye lens to present a text document, as illustrated in Figure 5a (left side). Most of the document is shown at a very small font, which gives the person a sense of the document's global structure. The user views local detail by selecting a focal point within the document, either by clicking the mouse on a line of text or by moving the scrollbar. If the scrollbar is used, the effect is that of sliding an optical lens up and down over the document. In Figure 5a, the user has selected line 157 as the focal point, and this line is shown in a large font. The surrounding 20 lines gradually decrease in size until the default background size is reached. Several hundred lines in total are visible here.

Users can tailor the shape and the magnification of their fisheye lens with the control panel shown on the right side of Figure 5a. First, they can adjust the font size of the background text or have it removed entirely. Second, users can change the shape of the lens that magnifies text around their focal point, using a custom-built lens widget. The black area of this widget represents a cross-section of the lens; users change the magnification function by moving any of the curve's points rightwards (to increase magnification) or leftwards (to decrease it). The curve is constrained to be always convex and symmetrical. As the lens is manipulated, the magnification function is immediately applied to the document.

This fisheye text viewer is also a groupware system that lets multiple people view the same document. People can join into this conference through GroupKit's session manager (Roseman and Greenberg, 1995)—shown on the right of Figure 5b—and this particular session (called 'Fisheye New') has three people joined into it. Within the fisheye viewer, the same document is presented on all displays. If a person loads a different document (an option on the File pull-down menu), all participants will see immediately the new text. As with the graphical browser, views are relaxed-WYSIWIS, for each person can set their own focal point and therefore have different fisheye effects.

Support for workspace awareness involves representing each participant's focus in the document. As in the graphical fisheye browser, location information is

[1] For simplicity, this prototype does not show node contents.

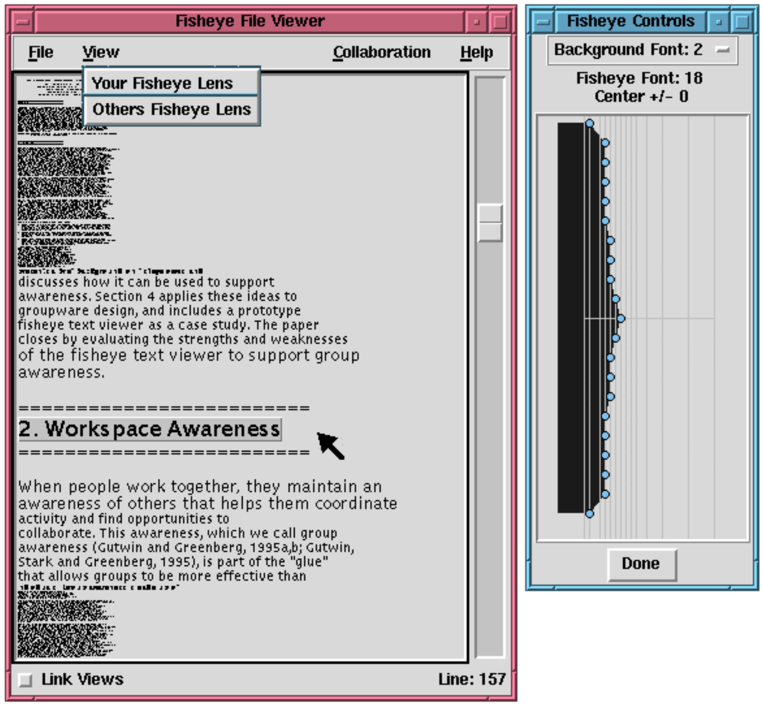

a) the single user version of the Fisheye Text Viewer

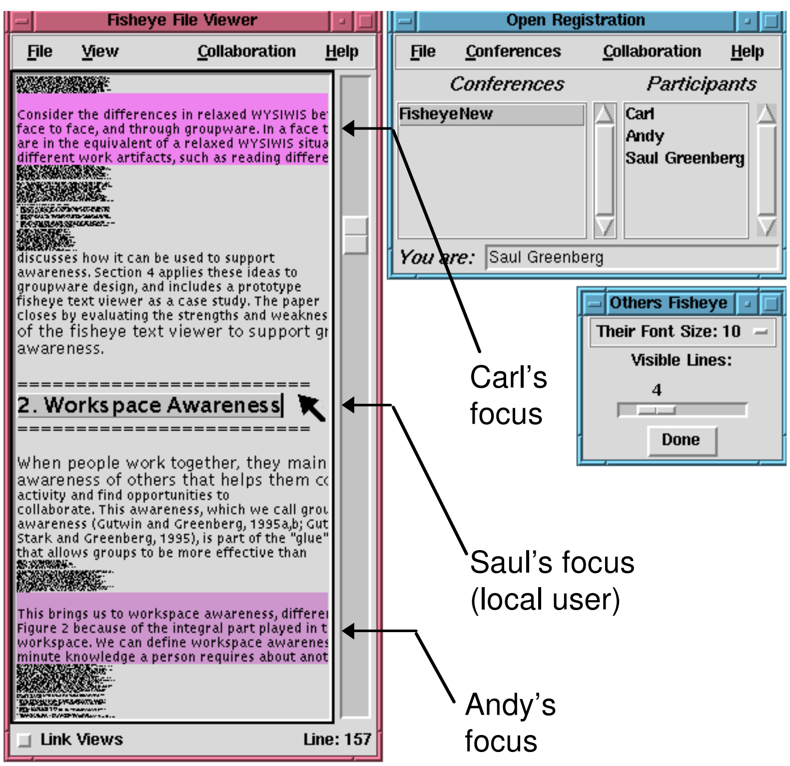

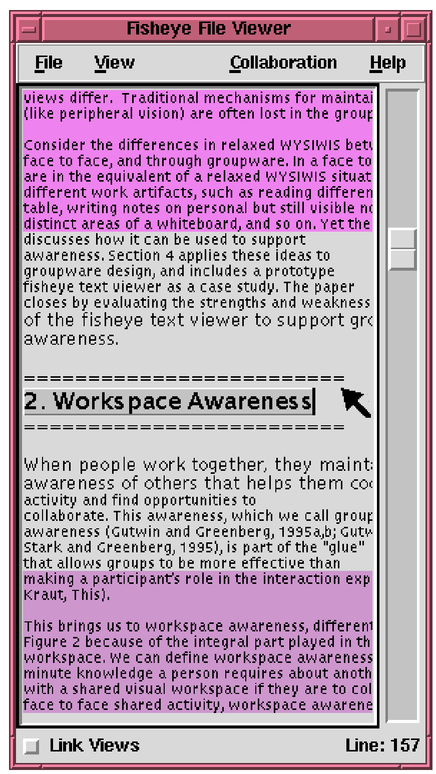

b) groupware fisheye with multiple focal points and global context

c) removing global context

Figure 5. The groupware fisheye text viewer

presented by marking others' focal points with their chosen colour. In addition, the text around other participants' focal points is also magnified. Figure 5b illustrates this: there are three focal points with corresponding magnified regions, the center region belonging to the local user and the surrounding two representing the other participants. Their locations in the global context and the detail of their work are clearly visible.

A user can also change the magnification function applied to other people's focal points—albeit in a simpler fashion—via the control panel on the middle right of Figure 5b. Moving the slider adjusts the range of the magnified region (here, to four lines), and a menu allows the font size of that region to be set (here set to 10 point font).

These fisheye controls allow users to flexibly allocate screen space for their own work or for the display of awareness information, as their tasks require.

- If only location information is required, a user can turn off the magnification of other participants' focus points. Their location will still be indicated through colour, but no detail will be shown. No extra screen space is used.
- When finer-grained awareness is desired, both location and detail can be progressively controlled by increasing the magnification around the other participant's focus, as well as the extent of the region being magnified.
- When people are working far apart in the document, a "split window" effect can bring them closer together. This is achieved by making the global view invisible (by displaying it at 0 point), thus showing only the magnified regions surrounding each focal point. For example, Figure 5c shows a split windows effect using the same focal points and document seen in Figure 5b.
- For tightly-coupled collaboration, people can align their views to something closer to a strict WYSIWIS situation in two ways. First, moving their own focal point onto another person's focal point is appropriate for quick and spontaneous interaction. Second, people can link their views by selecting the "Link Views" check button on the bottom left. This option lets all participants share a common focal point; if any user changes the focus, it will be changed on all other displays as well. View linking supports longer and more tightly coupled collaborations.

The fisheye text viewer has also been modified to cluster location information based on the document's semantic structure. For example, a code-viewing application places remote focal points on the name of the subroutine where that person is working, instead of showing the actual (and perhaps less meaningful) line of code.

4. Discussion

In face-to-face situations, people maintain a sense of awareness through practices that use the natural affordances of shared physical workspaces. Many of these natural affordances are lost in the transition to a groupware workspace, and groupware designers must explicitly build new mechanisms into their interfaces. The two prototypes just presented use the affordances offered by fisheye views as a way to better support existing mechanisms for staying aware of others in a workspace.

There is no guarantee, however, that a designer's new mechanisms will be appropriate for a particular groupware situation, and so the fit and the effectiveness of inventions like the groupware fisheye views must be evaluated. In order to structure thinking about support for workspace awareness, and to provide a vocabulary for classifying and comparing designs, we have begun our evaluation efforts by constructing a conceptual framework of workspace awareness (Gutwin, Greenberg and Roseman 1996).

This framework, summarized by the first column of Table 1, divides workspace awareness into several elements. The column lists several elements of workspace awareness, and lists questions that indicate what information a groupware system should capture and present to the group. These elements allow a type of heuristic evaluation (Nielsen 1993) that identifies whether a particular widget supports each element, and how the support is achieved. We have completed an initial heuristic evaluation of the two fisheye views discussed above. These results are summarized in the remaining columns of Table 1, and show strengths and weaknesses of the widgets' support for each element.

As the table illustrates, the two fisheye views provide a variety of information that covers several elements of workspace awareness. The main strengths of the approach, however, are in conveying information about location ('where') and about activity ('what'). Both widgets show other participants' locations within the workspace by highlighting and colouring their focal

Awareness Element	Support in the graphical fisheye browser	Support in the fisheye text viewer
Who • Who is participating?	✔Other participants are represented by colour ✘A coloured region may overlap and hide another ✘Remembering the colours assigned to individuals requires some effort	✔Other participants are represented by colour and magnified regions ✘A coloured/magnified region may overlap and hide another ✘Remembering the colours assigned to individuals requires some effort ✘A region may be out of view
What • What are they doing? • What tools are they using? • What are their intentions? • What changes are they making?	✔Others' pointers are shown and mapped appropriately ✘Actions are not shown in detail unless participants have the same focus	✔Area around each person's focal point is enlarged; details are clearly visible ✘Text cursors are (currently) not shown ✘A person's focal point may be scrolled out of view
Where • Where are they working? • What can they see? • Where are they pointing?	✔Others' focal points are indicated with node outlines ✔Others' focus areas are roughly shown with "halos" ✔Others' pointers are shown and mapped appropriately	✔Others' focal points are indicated with coloured lines and magnification ✔Focal point can act as cursor ✘Enlarged areas do not indicate actual viewport sizes ✘A person's focus point may be scrolled out of view
When • When have changes been made?	✔Changes are shown as they are made ✘Details of others' changes are only shown if participants have the same focus ✘No ability to replay past events	✔Changes are shown as they are made ✔Fine grained details are shown in the enlarged area ✘The user may miss changes in the global view when attending to the local view ✘No ability to replay past events

Table 1. Using elements of awareness for a heuristic evaluation of the two fisheye viewers

points. In addition, the graphical fisheye browser shows the rough extents of another person's focus area with a halo effect around the focus point (this capability could easily be added to the text viewer). By showing each person's focus area, the widgets give a rough idea of where people are likely to make changes, and which objects they are likely to use in the future. The location information supplied by the widgets is integrated within a person's normal view of the workspace, and so it is more readily available to someone as they work on their own parts of the group task.

Awareness of activity ('what') is also supported. By implementing multiple focal points, the fisheye text viewer is able to show details of what is happening in each person's focus. Used as a groupware viewer, participants can see details of the text that others are reading. Since the current system is also a rudimentary groupware editor, this means that fine-grained typing actions made by other people can be observed. In addition, the text viewer's tailorable lenses allow users to make their own decisions about allocating screen space, letting them trade awareness information for screen space and greater individual focus when their tasks require it.

However, some awareness elements are poorly supported. First, using colour to represent presence in the workspace is problematic, since it requires that people maintain a mapping between colours and individuals. Determining who is represented by a particular colour becomes difficult if there are more than a few people in the workspace. Second, awareness of past changes is not well supported by our fisheye views, although amount and recency of change could be factored into calculations of an object's interest value (Furnas 1986). Third, since a large document may not fit within the text viewer's main window, some focal points may be out of view. This problem arises because the current implementation requires at least one line of pixels for every line of text, and could be alleviated by using conventional fisheye techniques such as clustering, filtering, or true graphics scaling instead of fixed size fonts (Schaffer, 1995).

There are several ways that these prototypes can be improved to better support the maintenance of workspace awareness. Possible improvements include:

- addition of multiple focal points to the graphic fisheye browser;
- replacement of the text viewer's simple controls in Figure 5b with the flexible control shown in Figure 5a;
- allowing users to control magnification parameters for each remote participant;
- allowing users to specify multiple focal points if they are interested in several different parts of the document.

While heuristic evaluations do not take the place of user studies, they are quite appropriate to catalog expected advantages and disadvantages in design prototypes, as we have done here. The next step is to repair some, if not most, of the deficiencies mentioned in the Table. Many just require minor enhancements to the system e.g., we have already added multiple cursors to the text fisheye. Other enhancements will likely demand quite different techniques to support particular awareness needs, such as replay of past events for reviewing 'when' something had happened. When the obvious problems are solved, we can turn to user-based studies. This could include a comparison of fisheye views and traditional radar views, which would examine differences in people's abilities to interpret awareness information presented in these two formats.

5. Conclusions

In this paper, we have described how fisheye views of shared workspaces can be used as a basis for supporting workspace awareness in groupware. We presented two novel[2] prototype groupware systems that illustrate the strengths of fisheye views for showing awareness information. First, fisheye views provide a seamless transition between local detail and global overview, so that information about other participant's locations can be integrated within the normal view of the workspace. Second, fisheye views can contain multiple focal points, which can be used to show details of each participants' actions.

We also heuristically evaluated these fisheye prototypes using an initial awareness classification scheme. We identified their strengths and weaknesses, which will serve as the basis for modification and eventually a more focused evaluation of the effectiveness of applying fisheye views to groupware. Our experiences indicate that fisheye views hold promise as a means for helping people maintain workspace awareness, and by doing so, can improve the usability of real-time distributed groupware systems.

Availability. Both systems were implemented as shown in GroupKit, a groupware toolkit. GroupKit is available via anonymous ftp. The release contains all the software, installation instructions, example conference applications and session managers, manual pages, and tutorial documentation. The fisheye text viewer is part of the distribution, while the graphical fisheye is available from the author.

site:	ftp.cpsc.ucalgary.ca
directory:	pub/projects/grouplab/software.
http:	http://www.cpsc.ucalgary.ca/projects/grouplab/
mail list:	groupkit-users@cpsc.ucalgary.ca

Acknowledgements. This research is (gratefully) supported by the Intel Corporation and the National Engineering and Research Council of Canada. Neville Churcher, Ying Leung, Mark Apperley and Mark Roseman contributed to the ideas in this paper:

References

Baecker R. (1993) *Readings in Groupware and Computer-Supported Cooperative Work: Assisting Human-Human Collaboration*, Morgan-Kaufmann, San Mateo, CA.

Baecker R., Glass G., Mitchell A., and Posner I. (1994) SASSE: The Collaborative Editor. *Proc. ACM* Conference *on Human Factors in Computing Systems*, Volume 2, 459-460.

Beaudouin-Lafon M. and Karsenty A. (1992) Transparency and Awareness in a Real-Time Groupware System. *Proc. ACM Symposium on User Interface Software and Technology*, 171-180.

Bly, S. (1993) Media Spaces. In Baecker, R. (ed.) *Readings in Groupware and Computer-Supported Cooperative Work: Assisting Human-Human Collaboration.* Morgan Kaufmann, San Mateo, CA.

Dourish P. and Bellotti V. (1992) Awareness and Coordination in Shared Workspaces. *Proc. ACM Conference on Computer-Supported Cooperative Work*, 107-114.

Egan D., Remde J., Landauer T., Lochbaum C. and Gomez L. (1989) Behavioral Evaluation and Analysis of a Hypertext Browser. *Proc. ACM Conference on Human Factors in Computing Systems*, 205-210.

[2] As far as we know, this is the first application of fisheye techniques to groupware support. The closest related work is by King and Leung (1994), who created a multi-user folding editor to expand and contract hierarchical views.

Furnas, G. (1986) “Generalized Fisheye Views.” *Proc. ACM Conference on Human Factors in Computing Systems*, 16-23, April, Boston, Mass., ACM Press.

Greenberg S., Hayne S., and Rada R. (eds.) (1995) *Groupware for Real-Time Drawing: A Designer's Guide*, McGraw-Hill Europe, Berkshire, England.

Gutwin C. and Greenberg S. (1995) Support for Group Awareness in Real-time Desktop Conferences. *Proc. 2nd New Zealand Computer Science Research Students' Conference*, University of Waikato, Hamilton, New Zealand, April 18-21. Also as Report 95/549/01, Dept of Computer Science, University of Calgary, Calgary, Canada.

Gutwin C., Greenberg S. and Roseman M. (1996) Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. *Proc British Computer Society HCI ’96*, August 20-23, Cambridge University Press, to appear.

Gutwin C., Roseman M. and Greenberg S. (1996) A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. Report 96/585/05, Dept of Computer Science, University of Calgary, Calgary, Canada.

Gutwin C., Stark G., and Greenberg S. (1995). Support for Group Awareness in Educational Groupware. *Proc. Conference on Computer Supported Collaborative Learning*, Bloomington, Indiana, October 17-20, Distributed by LEA Press.

Heath C. and Luff P. (1991) Collaborative Activity and Technological Design: Task Coordination in London Underground Control Rooms. *Proc. European Conference on Computer Supported Cooperative Work*, 65-80, Amsterdam.

Ishii H. and Kobayashi M. (1992). ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact. *Proc. ACM Conference on Human Factors in Computing Systems*, 525-532, 705-706.

King R. J. and Leung Y. K. (1994). Designing a user interface for folding editors to support collaborative work. *Proc. British HCI Conference*. In People and Computers, #9.

Leland M., Fish R. and Kraut R. (1988). Collaborative Document Production Using Quilt. *Proc. ACM Conference on Computer-Supported Cooperative Work,* 206-215, Portland, Oregon, September 26-28.

Nielsen, J. (1993) *Usability Engineering*. Academic Press.

Roseman M. and Greenberg S. (1996) Building Real Time Groupware with GroupKit, a Groupware Toolkit. *ACM Transactions on Computer Human Interaction*. In press.

Sarkar M. and Brown M. (1992) Graphical Fisheye Views of Graphs. *Proc. ACM Conference on Human Factors in Computing Systems*, 83-91.

Sarkar M., Snibbe S., Tversky O. and Reiss S. (1993) Stretching the Rubber Sheet: A Metaphor for Visualizing Large Layouts on Small Screens. *Proc. ACM Symposium on User Interface Software and Technology*, 81-91.

Schaffer D. (1995) Visualizing Large, Loosely-Structured, Hierarchical Information Spaces. M.Sc. thesis, Dept of Computer Science, University of Calgary, Calgary, Canada.

Schaffer D., Zuo Z., Greenberg S., Bartram L., Dill J., Dubs S. and Roseman M. (1996) Navigating Hierarchically Clustered Networks through Fisheye and Full-Zoom Methods. *ACM Transactions on Computer Human Interaction*. In press.

Segal L. (1994) Actions Speak Louder Than Words: How Pilots Use Nonverbal Information for Crew Communications. *Proc. Human Factors*, 21-25.

Smith R., O'Shea T., O'Malley C., Scanlon E. and Taylor J. (1989) Preliminary Experiences with a Distributed, Multi-Media, Problem Environment. *Proc. 1st European Conference on Computer Supported Cooperative Work*, Gatwick, U.K., September 13-15.

Stefik M., Bobrow D. G., Foster G., Lanning S. and Tatar D. (1987). WYSIWIS Revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Office Information Systems*, **5**(2), 147-167, April.

Tang J. (1991). Findings from Observational Studies of Collaborative Work. *International Journal of Man Machine Studies*, **34**(2), 143-160, February.

Tatar D., Foster G. and Bobrow D. (1991). Design for Conversation: Lessons from Cognoter. *International Journal of Man Machine Studies*, **34**(2), pp. 185-210, February.

Valacich J., Dennis A. and Nunamaker Jr J. (1991). Electronic Meeting Support: The GroupSystems Concept. *International Journal of Man Machine Studies*, **34**(2), pp. 262-282, February.