Using Distortion-Oriented Displays to Support Workspace Awareness

Saul Greenberg[1], Carl Gutwin[1] and Andy Cockburn[2]

[1]Department of Computer Science
University of Calgary
Calgary, Alberta
Canada T2N 1N4
phone: +1 403 220 6087
email: saul *or* gutwin@cpsc.ucalgary.ca

[2]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 moving away from strict view-sharing and towards relaxed "what-you-see-is-what-I-see" (relaxed-WYSIWIS) 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 propose *distortion oriented displays* as a novel way of providing this awareness. These displays, which employ magnification lenses and fisheye view techniques, show global context and local detail within a single window, providing both peripheral and detailed awareness of other participants' actions. Three prototypes are presented as examples of groupware distortion-oriented displays: the *fisheye text viewer,* the *offset lens*, and the *head-up lens.*

Keywords: awareness, magnifying lenses, fisheye views, distortion-oriented displays, desktop conferencing, groupware.

1. Introduction.

Real-time distributed groupware allows people who are geographically separate to work together at the same time through computers. These systems provide a shared virtual workspace where conference participants can see and manipulate work artifacts. The shared workspace typically contains groupware tools such as a shared sketchpad or drawing system (e.g., Greenberg, Roseman, Webster and Bohnet 1992), multi-user text editors (e.g., Baecker, Glass, Mitchell and Posner 1994), idea organizers (e.g., Tatar, Foster and Bobrow 1991) or multi-user games. In addition to the workspace, a groupware system is likely to incorporate facilities for communication over audio and video links.

Unfortunately, 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 where all participants see exactly the same view of the workspace at all times (Stefik, Bobrow, Foster, Lanning and Tatar 1987). The relaxations give people control over their own view of the 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, however, can contribute to a loss of awareness since, when views differ, people can lose track of where others are and what they are doing in the workspace. One technique to support awareness in relaxed-WYSIWIS provides users with two separate windows: a normal-sized view of one's own working area; and a "radar" overview that shows a miniature of the entire workspace, typically overlaid with boxes that represent each participant's viewport (e.g. Smith, O'Shea, O'Malley et. al. 1989; Baecker, Glass, Mitchell & Posner 1994). While these work well in some tasks (Gutwin, Greenberg and Roseman 1996), the separate windows introduce a physical seam between local and global contexts that a user may find difficult to integrate, and the radar miniature may not have enough resolution to show the necessary details of anothers' activity.

In this paper, we propose *distortion oriented* displays as a mechanism for presenting awareness information. These displays show both global context and local detail within a single window. They work by scaling most or all of a workspace to fit within a window, and then distorting (or magnifying) a region to show its detail. When applied to groupware, the distortion oriented display provides both peripheral and detailed awareness of other participants by showing their position and actions in the global context, and by distorting the area around their work to see the details of the interaction.

In the following subsection we briefly review the workspace awareness requirements that groupware systems should satisfy.

1.1. Workspace Awareness

When people work together, they maintain an awareness of others that helps them coordinate activity and find opportunities to collaborate (e.g. Dourish and Belloti 1992). This awareness, which we call *group awareness* (Gutwin, Stark and Greenberg 1995; Gutwin and Greenberg 1996), is part of the "glue" that allows groups to be more effective than individuals. Group awareness is made up of several kinds of knowledge about what is happening in one's collaborative environment, as summarised below.

Informal awareness of a work community is basic knowledge of who is around in general (but perhaps out of sight), who is physically in a room with you, and where people are located relative to you. *Group-structural awareness* involves knowledge about such things as people's roles and responsibilities, their positions on an issue, their status, and group processes. *Social awareness* is the information that a person maintains about others in a social or conversational context: things like whether another person is paying attention, their emotional state, or their level of interest. The fourth kind of group awareness is *workspace awareness*, which involves knowledge about how the others in the group interact with the shared workspace. In a face-to-face interaction, the shared workspace is often the tabletop and whiteboard, where people bring artifacts such as documents to the table, pass them to each other, point and gesture around them, use tools to modify them, and make notes on white-boards.

We define workspace awareness more precisely as the up-to-the minute knowledge a person requires about another group member's interaction with a shared workspace if they are to collaborate effectively. While it is less easy to define exactly what knowledge people require, the first column in Table 1 in Section 4.2 summarises a few of the more essential elements comprising awareness, phrased as questions(the framework is reviewed in Gutwin, Greenberg and Roseman 1996 (this proceedings) and more fully in Gutwin et al. 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 heuristic guidelines for the development of the awareness prototypes, as described in Section 3.

The following section presents a brief background on distortion oriented displays. Section 3 then introduces three prototypes that demonstrate the application of distortion-oriented views in groupware: the *fisheye text viewer,* the *offset lens*, and the *head-up lens*: a video is also available that illustrates their dynamics (Greenberg, Gutwin and Cockburn, 1996). The paper closes by discussing both the strengths and weaknesses of using distortion-oriented displays to support group awareness.

2. Distortion-Oriented Displays

A central concern in information visualisation is how a system can present both global structure (that provides overview and context) and local detail (that reveals information in the user's area of interest). Distortion-oriented displays allow visualisation that merges the global view of the information and the local detail of interest to the user. These displays can be categorised into two approaches: magnifying lenses, and fisheye views. Each are briefly discussed below.

2.1. Magnification Lenses

When a paper document contains details too small for people to read, they can use a magnifying lens to enlarge a portion of it. Similarly, a magnifying lens metaphor can be applied to computer displays. At its simplest, consider a large workspace that is scaled to fit within a single window. This provides the viewer with a sense of the global context but poor detail. When the viewer points to an area of interest on the display, a separate "lens" window containing a magnified view of that area appears on top of the original one.

Computer-based magnifying lenses surpass their physical metaphors. Stone, Fishkin and Bier (1994) introduce the Magic Lens: a movable filter that affects the appearance of objects viewed through it, in ways that go far beyond simple magnification. Aside from scaling, they have applied the Magic Lens to show a variety of information: different renderings of pictures; state information of objects that is normally hidden; additional structures such as grids; selective detail of a view; and so on. They also show how a lens can be turned into a *click-through tool*, which modifies a user's input over the transformed region being viewed. A taxonomy of such see-through tools is given by Bier, Stone, Fishkin, Buxton and Baudel (1994).

2.2. Fisheye Views

Fisheye views are computer visualisation techniques that provide both local detail and global context in a single display (this concept was originally coined a *bifocal display* by Spence and Apperley in 1982). Unlike magnification lens techniques, where entities are either magnified or not, fisheye views display the global context and local detail on a continuous "surface." The user chooses a point of focus where they wish to see local detail: this area is visually emphasised, and the remainder of the data is made less visually important.

Fisheye views have been used to visualise data in many domains. Furnas (1986) created systems for viewing and filtering structured program code, biological taxonomies, and calendars. Egan, Remde *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.

Although fisheye techniques are normally used to emphasise a single focus point, multiple focus points can also be supported. 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.

3. Applying Distortion-Oriented Views to Group Awareness

In this paper, we suggest that distortion-oriented views can also be applied to groupware as well as conventional information visualisation. We contend that distortion-oriented views are well-suited to groupware because of their ability to provide awareness of others' actions in the workspace within a single window. To achieve this awareness, the positions and coarse actions of all participants are displayed within the global context, while the magnified areas present the local details of each participant's particular interaction.

To illustrate how this can be done, we present three prototypes: the *fisheye text viewer*, the *offset lens*, and the *head-up lens.* This section describes the awareness features provided by the prototypes. The limitations of these systems, and our plans for further work, are discussed in Section 4.

3.1. Fisheye Text Viewer

The fisheye text viewer (Greenberg, Gutwin and Cockburn 1996) supports awareness by assigning one focal point to each participant, and by giving each person the ability to tailor the magnification function of any of the focal points. It reveals the location of others within its workspace, and illustrates how details of other people's activity can be presented via multiple focal points. To demonstrate the fisheye text viewer, we first present how it works as a single-user system, and then how it works as a multi-user system.

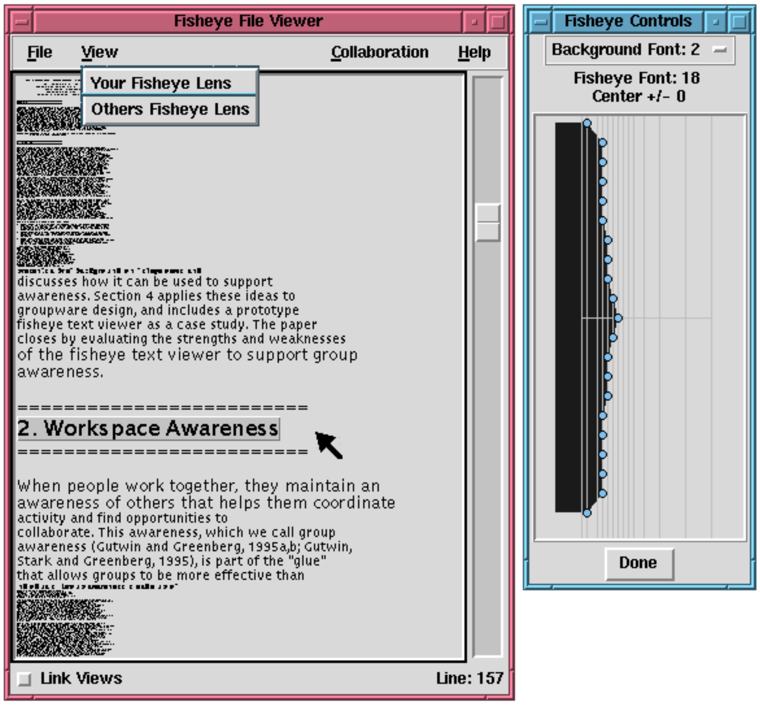

1a) the single user version of the Fisheye Text Viewer

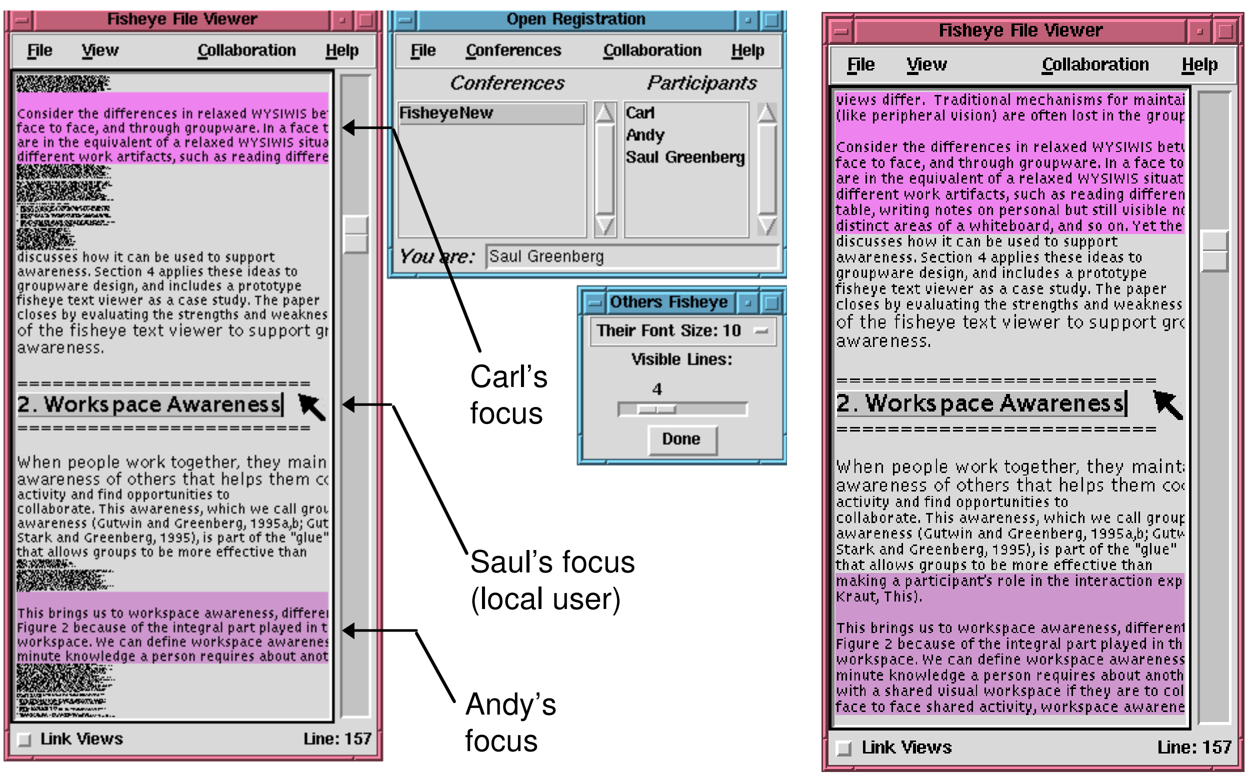

1b) groupware fisheye with multiple focal points and global context

1c) removing global context

Figure 1. The groupware fisheye text viewer

3.1.1. Single-User Fisheye.

The viewer uses a fisheye lens to present a text document, as illustrated in Figure 1a (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 1a, 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.

Users can tailor the shape and the magnification of their fisheye lens with the control panel shown on the right side of Figure 1a. First, they can adjust the font size of the background (global) 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 control. The black area of the control represents a cross-section of the lens; users increase the magnification function by moving any of the curve's points rightwards, 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.

3.1.2. Groupware Fisheye.

The fisheye text viewer is also a groupware system that lets several people view the same document. Each person's view is relaxed-WYSIWIS, allowing each individual to set their own focal point on the document.

Workspace awareness is supported by representing each participant's focus in the document. Referring to the awareness factors in the first column of Table 1, identity and location information are presented by marking others' focal points with their chosen colour. In addition, the text around other participants' focal points is also magnified. Thus, activity awareness is provided through each user's view of the other participants' focal points. Figure 1b illustrates this: there are three focal points with corresponding magnified regions, the centre region belonging to the user and the surrounding two representing the other participants. Their locations in the global context and the detail of their work are clearly, and immediately, visible to the other participants.

A user can also change the magnification function applied to their view of other people's focal points—albeit in a simpler fashion—via the control panel on the middle right of Figure 1b. 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, thus displaying only the regions surrounding each focal point. For example, Figure 1c shows a split windows effect using the same focal points and document seen in Figure 1b.
- 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 clicking a check button on the viewer (Figure 1b+c, 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.

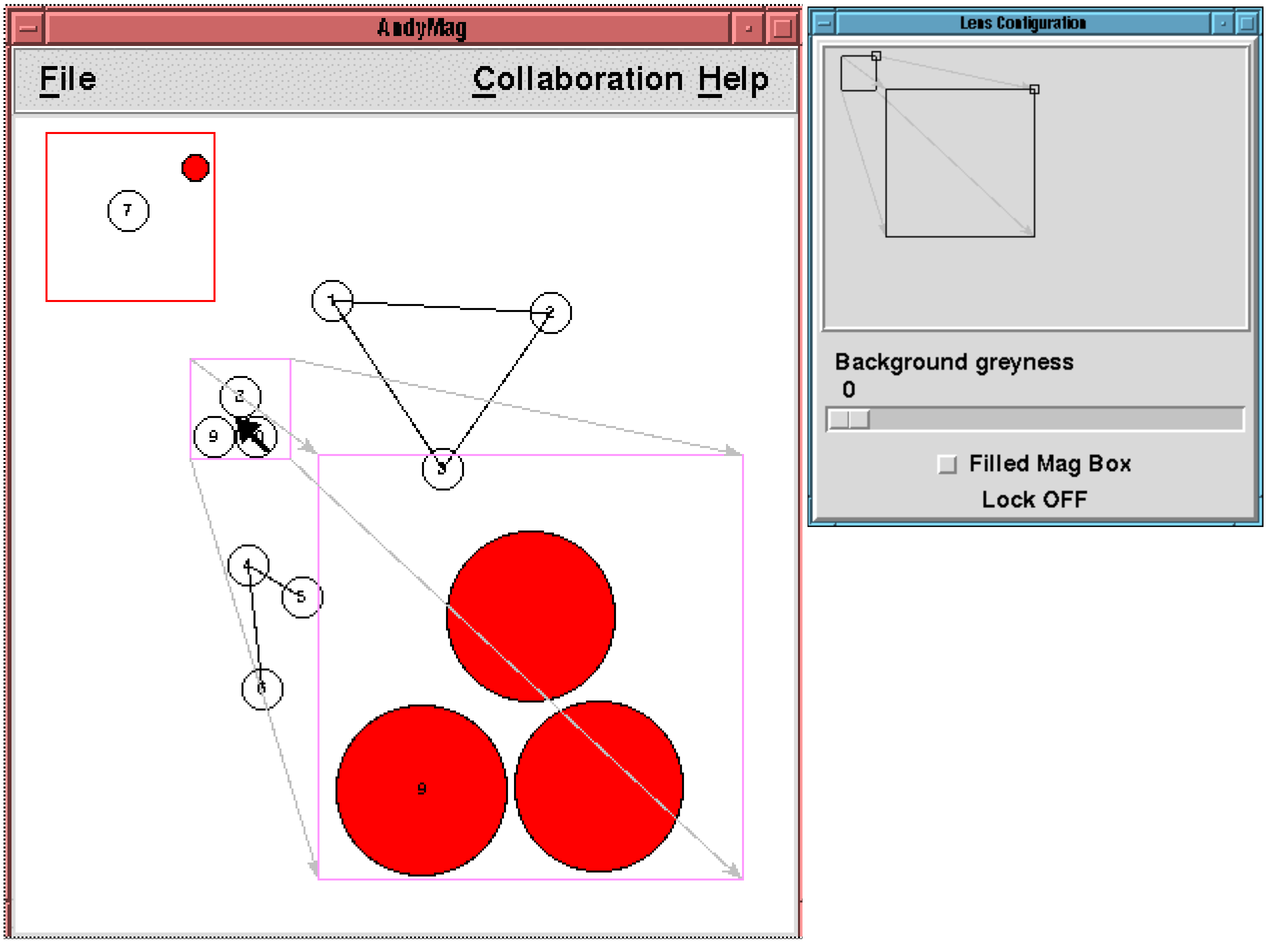

Figure 2: The Offset Lens, showing the magnification and offset controls on the right

3.2. The Offset Lens

The Offset Lens is a magnification-oriented system that allows participants to view and concurrently edit a shared graphical workspace. Figure 2 shows the Offset Lens in use. Nodes in the graph represent arbitrary 'grains' of information—for instance, pages in a hypertext document, a concept map, or independent design decisions and their associated rationale. The workspace is scaled so that the entire graph fits the window, and thus contains the entire 'global context.' In this section, the offset lens is first explained as a single-user system, and then as a groupware system.

3.2.1. Single-User Offset Lens.

The global context is directly editable, and the user can add new nodes and edges (lines) to it by clicking or dragging with the mouse. The user also has a magnification lens layered over a sub-area of the workspace, shown in Figure 2 as the bordered projected area on the bottom right. This lens shows the 'local detail' (the magnified nodes and lines around a person's cursor): the lens' position on the display follows the cursor as one moves around the global context, and the contents of the local detail continually update to show the new sub-area beneath the lens. As a single user system, it is similar in spirit to Ware and Lewis' (1995) DragMag image magnifier (the original Offset Lens, created before we saw the DragMag, was changed to include some of DragMag's single user features).

Like the global context, the local detail region is editable. Editing the local detail requires the user to lock the magnification lens' position onto the current focus (by clicking with the right-hand mouse button). Subsequent editing actions, identical to those at the global-level, take effect at the current magnification level. The cursor location and editing actions are immediately reflected, to all users, in the global context. The user can therefore see the consequences of edits in both the local detail and global contexts.

The user can alter several of the lens properties through the control panel (Figure 2, right window): the size of the lens; the position of the lens relative to the focal point; and its transparency.

1. The balance between global view and local detail is adjusted by directly setting the size of the magnifying lens. While the magnification factor remains constant, the area magnified by the lens (and thus the size of the lens) can be enlarged or made smaller.
2. A true magnifying lens, positioned over an unmagnified region, occludes some of the unmagnified nodes beneath the lens. This tends to hide the connections between the local detail and global context. To mitigate this problem, the system allows a user to offset the magnified image away from the focus by dragging the magnifier to a new position (Figure 2, right window). When the offset is zero, the magnified image is placed on top of the focus. Adjusting the amount of offset moves the lens from the focal point, eventually allowing simultaneous views of the area of interest in both global and local contexts.
3. The global context and local detail are drawn as two elements that can compete for attention, and a user may want to adjust their visibility to suit the task on hand. The Offset Lens allows this adjustment in two ways. First, the magnifying lens can be set to transparent or opaque, set by selecting the "Filled Mag Box". For instance, the task of browsing the work surface is easier if one can see through the magnifying lens, while focusing on local detail is easier if the lens is opaque. Second, the visibility of the global context can be controlled by having it drawn at different grey levels. By adjusting a slider, the greyness moves from black to white and the global context becomes less visible and the local detail stands out further.

3.2.2. Groupware Offset Lens.

The groupware Offset Lens uses relaxed WYSIWIS. All users see the same representation of the global context, including immediate updates following editing actions. However, each user's local detail (the images within their magnifying lenses) is not shared, allowing participants to focus on whatever part of the workspace they wish to view. Telepointers are also supported for gesturing and location awareness.

Referring to the awareness factors in Table 1, location awareness is provided by revealing to the other participants the position of each user's magnification lens on the global context. The identity of each user is determined by uniquely colouring their lens (another user is shown by the rectangle in Figure 2, top left). Activity and temporal awareness are provided by the immediate updates of editing actions on the global context. Additional location and activity awareness is available when a user is editing within their local detail (rather than in the global context). This information is communicated by small telepointers on the global context which provide an indication that the user is focused on a specific region and intends to carry out editing actions.

Through this combination of awareness mechanisms, each user can monitor the global context and stay aware of their colleagues' presence, their region of activity, where they are currently pointing to, and what actions they are doing. By not showing everyone's magnified views, a person's display is left uncluttered. Of course, people can align their magnified views when sharing of detailed information is required.

3.3. Head-Up Lens

The Offset Lens takes the local and global views of a workspace, and merges them into a single display that shows both views at the same time. The Head-Up lens, which is a graph editor, takes this one step further, by layering and resizing both the views to fit the window exactly. It is a "transparent layered user interface," as defined by Harrison, Ishii, Vicente and Buxton (1995).

3.3.1. Single-User Head-Up Lens.

As with most head-up displays, our lens provides a two-level view of the workspace. It is illustrated by the graphical editor in Figure 3. Like the Offset Lens, the global context shows the entire workspace, scaled to fit the size of the window exactly. The foreground shows the local detail which is a viewport onto a sub-area of the background global context. The location of the user's viewport onto the global workspace is controlled through scrollbars and other conventional interface mechanisms. The two primary differences between the Head-Up Lens and the Offset Lens are that, first, the Head-Up interface is simpler because there is no need to raise or position the lens, and second, the user is unable to edit the global context directly in the Head-Up system.

3.3.2. Groupware Head-Up Lens.

As with the Offset Lens, uniquely coloured rectangles on the global context show the view extents of the local and remote participants, providing location and identify awareness. For example, the foreground viewport in Figure 3 is

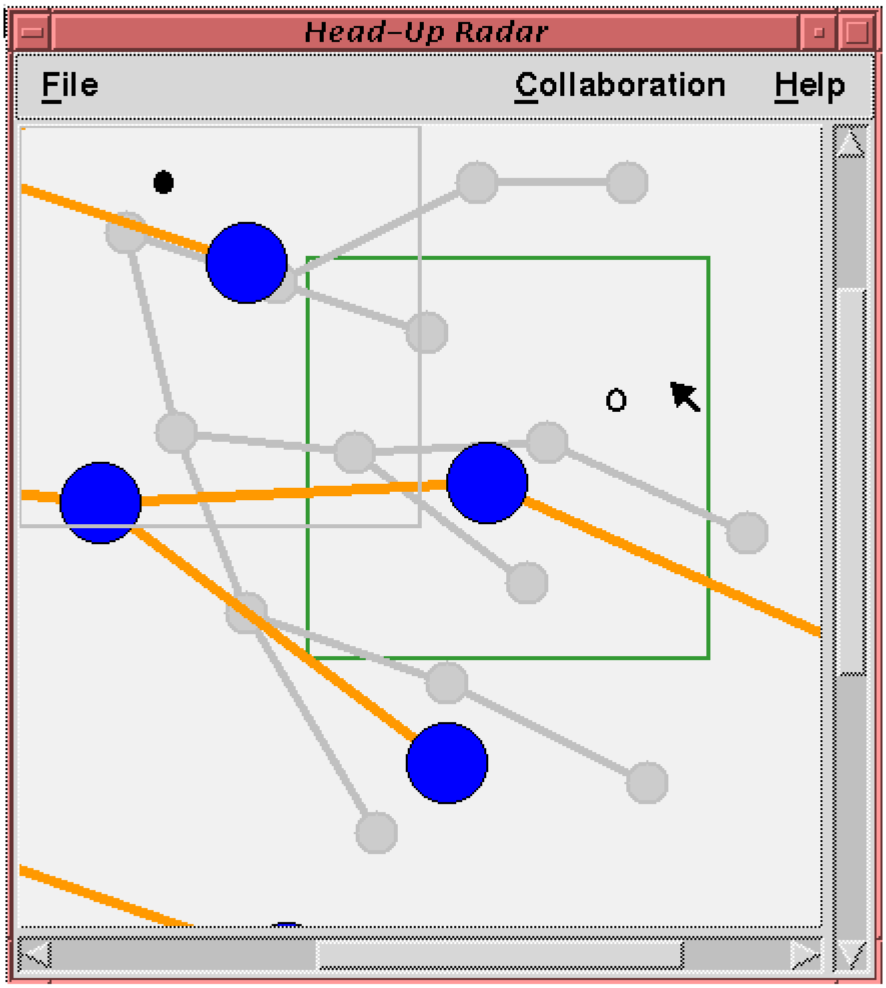

Figure 3: The Head-Up Lens.

reflected by the middle-right rectangle in the background. When someone moves their foreground view, their rectangle slides around the background, showing where they are currently located. In addition, miniature telepointers in the background (drawn as small circles) give some indication of what object others are focusing on, proving activity awareness. The telepointers on the global context in both the Offset and Head-Up Lens systems also allow limited gestural communication even when participants do not share the same local view. To reduce the amount that activities on the global context intrude on a person's attention on the local-detail objects, the background is "ghosted out" in light grey.

4. Discussion and Further Work

The previous section described the interfaces and features of three prototypes for providing workspace awareness in groupware. The current implementations are intended to be point systems indicating what is possible and emphasising the technical feasibility of group-awareness in distortion-oriented displays. In this section we identify the limitations of the prototypes, focusing on inadequacies in their interfaces and on mismatches between the awareness facilities they provide and those identified as desirable in Section 2. This discussion serves as a specification of our further work as we iterate from point-systems to evaluable working prototypes.

4.1. Assessing the User Interfaces

Although it is premature to run formal usability studies on the prototypes, the design team has experimented with each of the system. Our aim in assessing the interfaces is to remove the large grain usability flaws to ensure that subsequent usability analysis identifies problems with the support for workspace awareness rather than symptoms of lower-level interface errors.

Of the three prototypes, the fisheye-text viewer has the most polished interface. Users had few problems in changing their focal points and tailoring the focal properties of the lenses. The primary limitation of the fisheye prototype is functional, in that its editing facilities are rudimentary.

The most fundamental problem of the Offset and Head-Up Lens systems is that users are required to mentally integrate the magnified and unmagnified planes of work. The fisheye viewer does not suffer this problem so severely because the magnified regions appear on a continuous plane. Finding effective ways to balance people's need to both focus and divide their attention on transparent layered interfaces is a research issue in its own right, as now being explored by Harrison, Ishii, Vicente and Buxton (1995).

More generally, the interface to the Offset Lens is complex when compared to the two other systems. This is primarily due to the large number of user customisable features, but in addition, special interface measures are required to let a user edit both the global and local regions. The user has to select a mode which locks the lens onto the display, and directs user input to the local-detail region. Unfortunately, the locked lens makes it difficult to interact with objects that lie just outside the magnified area. This problem is partially resolved by the small telepointer which shows the user's area of action on the global context. An alternative (unimplemented) solution to the locking problem is based on Bier *et al.*'s (1994) two-handed input techniques—one hand is used to move the lens over the display, while the other hand controls the mouse cursor and the interaction to the area either inside or outside the lens.

The interface to the Head-Up Lens is simpler, but more constrained, than that of the Offset Lens. The locking problems are resolved because editing is only possible on the local detail layer. This limitation would be straightforward to remove by adding a toggle that flips global and local layers, and by redirecting the input to the global level. This modification would, however, come at the cost of additional interface complexity. Other powerful, but complex, controls in the Offset Lens are not available in the Head-Up Lens. These include controls for the lens' size, magnification function, and the degree of shading used to obscure the global-context. Despite its interface simplicity, the Head-Up Lens suffers problems similar to the Offset Lens, as well as a few others. Of particular note is the problem that changes in the global context, caused by the actions of others', can interfere and annoy a person concentrating on their local detail.

Some of the potential problems described above are repairable, others are ingrained in the fundamental approach of the particular distortion oriented technique. The ultimate viability of the systems, and the degree to which these *potential* problems affect users, has yet to be determined through user testing

4.2. Assessing the Workspace Awareness

Generally, the systems satisfy the criteria for location and activity awareness (where and what) more successfully than the criteria for identity and temporal awareness. In assessing the user interfaces of the prototypes, issues of awareness have already been raised: for instance, the fact that others' actions can impinge on a user's local detail in the Head-Up Display. There are many other trade-offs and problems in the awareness mechanisms supported by the three prototypes, as summarised in Table 1. Each of the forms of awareness is briefly discussed below.

4.2.1. Identity Awareness.

All the systems use colour coding as the main method of identifying participants. Each user, therefore, has a cognitive burden of mapping from colours to individuals. In our experience, this is not a problem as the small size of the group and the natural verbal and gestural deixis between participants strongly reinforce the colouring identification scheme. However, mapping could be difficult if the group is large or meets infrequently, or if speech channels are not immediately available. Another problem in the use of colouring occurs when overlapping colours obscure each other. This problem affects the fisheye text viewer most seriously, as large blocks of text may be coloured: in the two lens systems, only the bounding boxes of the lens regions are coloured. A partial solution to this problem would be to allow mouse actions within a region to pop up the names of those currently working in the area.

<table>
<tr><th>Awareness Element</th><th>Fisheye Text</th><th>Offset Lens</th><th>Head-Up Lens</th></tr>
<tr><td rowspan="2">Identity
Who is participating?</td><td rowspan="2">✔Visible as coloured region and as enlarged font
✘ One coloured region may overlap another.
✘ Area may be out of view</td><td colspan="2">✔Coloured viewport and cursors of others visible in global context
✘May be occluded by magnified objects</td></tr>
<tr><td>✔Lens may be offset to see what is below it</td><td>✘Foreground image must be scrolled to a clear area to make the occluded background visible</td></tr>
<tr><td>Location
Where are they working?
What can they see?
What are they pointing at?
Where can they have effects?</td><td>✔Focal point shown as coloured lines within global context
✔Area around the focal point enlarged
✔Focal point can act as cursor
✘Enlarged area does not represent actual viewport size</td><td colspan="2">✔All viewport shown within global context as coloured boxes
✔Small cursors shown in global view
✘Images within viewport may be too small to determine what a person can see</td></tr>
<tr><td>Activity
What are they doing?
What are their intentions?
What changes are they making?</td><td>✔Area around the focal point enlarged, with details clearly visible
✘Text cursor not shown
✘Area may be scrolled out of view</td><td colspan="2">✔Changes made in detailed view immediately visible in global view
✔Fine-grained movement of small cursors in global view indicates intent
✘Global view may not have enough detail to make changes and cursor movement comprehensible</td></tr>
<tr><td>Temporal immediacy
When have changes been made?</td><td colspan="3">✔Changes are shown as they are made
✘No ability to replay past events
✘Can miss changes in the global view when attending the local view</td></tr>
</table>

Table 1. Assessing the prototypes with respect to awareness criteria.

4.2.2. Location Awareness.

Extensive use of telepointers and moving viewports provide rich information on the region of participants' activity in each of the systems. Because all participants' locations are embedded in the global view, it is easy for a user to situate exactly where others are working.

Awareness affords opportunities for tightly-coupled interaction, and as a consequence the ability to couple locations (and therefore views) would be useful. Currently, only the fisheye text viewer provides an explicit facility for tightly coupled views of the workspace. Thus, in the Offset and Head-Up Lens systems, users who wish to work directly on the same section of the workspace must make the necessary view adjustments independently to ensure that their focal regions are similar. Future versions of the Offset and Head-Up Lens systems will provide a view linking option, similar to that in the fisheye text viewer.

4.2.3. Activity Awareness.

By implementing multiple focal points, the fisheye text viewer is able to show details of what is happening in each person's focus. 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, the region of other participants' activity may be scrolled out of view if the document is large. In contrast, in the Offset and Head-Up Lens systems, the scrolling problem does not occur because the global context reveals the entire workspace at all times. The lens systems are, however, susceptible to another problem—in a very large workspace, the global context may lack the detail to provide useful activity awareness.

The counter-side to activity awareness is clutter. When focusing on the details of their personal work, users are likely to want a dedicated view that masks background activity. The tailorable lens in the fisheye text viewer allows the user to suppress information about the activity of others, and the Offset Lens allows the user to mask out the global context. The Head-Up Lens, as currently implemented, makes no user configurable allowance for the suppression of activity information, but this could be easily repaired at the cost of additional interface complexity.

4.2.4. Temporal Awareness.

Although all show updates as soon as they are made, none of them support awareness over a period of time. If a user leaves the session for a period, or misses a sequence of updates because the region was obscured or scrolled out of view, there is no support for finding out what has changed, for replaying the sequence of actions, or for finding out who did what.

4.3. Summary

Assessing the prototypes' support with respect to the awareness criteria is useful in helping us identify potential problems prior to end-user evaluation. What is clear is that the distortion-oriented techniques do, at least in theory, support many awareness needs. Of course, there is no guarantee that users can use this information in practice. The benefits and problems that emerge in actual use are yet to be determined in usability studies.

5. Conclusions

In this paper, we have identified the lack of workspace awareness as a major limitation in current relaxed WYSIWIS groupware. The critical factors in workspace awareness were discussed, and distortion-oriented visualisation techniques were proposed as a technology for satisfying many awareness requirements. Distortion-oriented techniques are promising because they allow awareness information to be integrated within large information spaces, while minimising the demands on screen real-estate.

The three prototype groupware applications described in the paper demonstrate novel ways that distortion-oriented displays can provide people with a sense of group awareness. The capabilities of these systems were assessed with respect to the workspace awareness criteria. While much work remains to be done, we believe that the awareness facilities demonstrated by these systems will ultimately improve the usability of real-time distributed groupware.

These distortion-oriented awareness tools are all derived from single-user equivalents. We believe that these techniques will be *at least* as useful as their single-user counterparts, for the groupware extensions make no constraints on single-user use. We also believe that leveraging these techniques to support group work will make them even more beneficial.

Availability.

GroupKit, the toolkit used to implement the awareness prototypes, is available via anonymous ftp. The actual systems described in this paper are either included in the release, or available from the authors.

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

Acknowledgements.

This research is (gratefully) supported in part by Intel Corporation, and the National Engineering and Research Council of Canada. Neville Churcher and Mark Roseman contributed in one way or another to this work.

References

Baecker R., Glass G., Mitchell A., and Posner I. (1994). SASSE: The Collaborative Editor. *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, Volume 2, pp. 459-460.

Bier, E.A., Stone, M.C., Fishkin, K., Buxton, W., and Baudel, T. (1994) A taxonomy of see-through tools. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, p358-364, April 24-28, Boston, Mass., USA, ACM Press.

Dourish P. and Bellotti V. (1992). Awareness and Coordination in Shared Workspaces. In *Proceedings of ACM CSCW'92 Conference on Computer-Supported Cooperative Work*, pp. 107-114.

Egan D. E., Remde J. R., Landauer T. K., Lochbaum C. C., and Gomez L. M. (1989). Behavioral Evaluation and Analysis of a Hypertext Browser. *Proceedings of ACM CHI'89 Conference on Human Factors in Computing Systems*, pp. 205-210.

Furnas, G. (1986) Generalized fisheye views. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, p16-23, April, Boston, Mass., USA, ACM Press.

Greenberg S., Gutwin, C. and Cockburn, A. (1996a). Awareness Through Fisheye Views in Relaxed-WYSIWIS Groupware. *Proceedings of Graphics Interface*, p28-38, Toronto, Canada, May 21-24. Distributed by Morgan-Kaufmann.

Greenberg, S., Gutwin, C. and Cockburn, A. (1996b) Applying distortion-oriented displays to groupware. Videotape available from the authors, Department of Computer Science, University of Calgary, Alberta, Canada.

Greenberg, S., Roseman, M., Webster, D. and Bohnet, R. (1992). Human and technical factors of distributed group drawing tools. *Interacting with Computers*, 4(1), p364-392. Butterworth-Heinemann.

Gutwin, C., Greenberg, S. and Roseman, M. (1996). Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. People and Computers XI (Proceedings of the HCI'96), A. Sasse, R.J. Cunningham, and R. Winder, Editors. Pages 281-298, Springer-Verlag. Conference held at Imperial College, London, August 20-23.

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

Harrison, B., Ishii, H., Vicente, K. and Buxton, W. (1995) Transparent layered user interfaces: An evaluation of a display design to enhance focused and divided attention. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, p317-324, May 7-11, Denver, Colorado, USA, ACM Press.

Sarkar M. and Brown M. H. (1992). Graphical Fisheye Views of Graphs. *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, pp. 83-91.

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

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 TOCHI Transactions on Computer-Human Interaction*, 3(2), p162-188, June.

Smith R. B., O'Shea T., O'Malley C., Scanlon E., and Taylor J. (1989). Preliminary Experiences with a Distributed, Multi-Media, Problem Environment. In *Proceedings of the 1st European Conference on Computer Supported Cooperative Work (EC-CSCW '89)*, Gatwick, U.K., September 13-15, Computer Sciences House, Sloug, UK.

Spence, R. and Apperley, M.D. (1982). Data-Base Navigation: An Office Environment for the Professional. *Behaviour and Information Technology*, **1**(1), pp. 43-54

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), pp. 147-167, April.

Stone, M.C., Fishkin, K. and Bier, E.A. (1994) The MovableFilter as a User Interface Tool." *Proceedings of the ACM SIGCHI* Conference *on Human Factors in Computing Systems*, p306-312, April 24-28, Boston, Mass., USA, ACM Press.

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

Ware, C. and Lewis, M. (1995). The DragMag Image Magnifier. *Chi '95 Video Program.* From ACM Conference on Human Factors in Computing System, May 7-11, Denver. ACM Press. Videotape.