

An Investigation of Groupware Support for Collaborative Awareness Through Distortion-Oriented Views

Andy Cockburn
Department of Computer Science
University of Canterbury
Christchurch, New Zealand
andy@cosc.canterbury.ac.nz

Philip Weir
Compaq Application Development Centre
Christchurch, New Zealand
Philip.Weir@Compaq.com

Contact author: Andy Cockburn

Abstract

This paper reviews models and theoretical frameworks of collaborative awareness in the use of real-time groupware systems. The review is used to motivate and guide an investigation of distortion-oriented mechanisms for supporting collaborators' fluid and dynamic awareness requirements. We describe our development and evaluation of DOME, a distortion-oriented multi-user editor. Although we designed DOME to provide a realistic and useful platform for the investigation of awareness concepts, our evaluation revealed major flaws in its support for distortion-oriented awareness. We analyse the cause of these errors, some of which have been undetected in prior work, and provide precise formulations that will overcome them.

Keywords:

Groupware, collaborative workspace awareness, distortion-oriented visualisation

1 Introduction

When working with others on physical media such as white-boards and desktops, collaborators maintain a continual awareness of their colleagues. This awareness is seamlessly managed through rich and varied channels such as glances “out of the corner of the eye” which do not disrupt each individual's activity. The subtlety of collaborators' use of awareness information is vividly illustrated in Heath and Luff's (1992) ethnographic study of workplace activity in a London Underground control room.

“... the Controller will swear, feign momentary illness or even sing a couple of bars of a song to draw the DIA's attention to an emergent problem within the operation of the service. The various objects used [...] allow the individual to continue with an activity in which they might be engaged, whilst simultaneously inviting them to carefully monitor a concurrent event”, page 81.

The provision of awareness within computer-supported media-spaces is a major problem for groupware developers. The simplest “solution” is to enforce a strict-WYSIWIS (what you see is what I see) constraint on the views of all users. The advantages of strict-WYSIWIS are that it is relatively easy to implement, and that it is conceptually easy to use because all users have precisely the same view. Its major limitation, however, is that it only allows a highly constraining style of collaboration in which all users are forced to work as a tightly coupled unit. This fails to support the natural and dynamic movement between different styles of collaboration in which, for instance, periods of close collaboration follow periods of autonomous work.

Early experiments with relaxing the strict-WYSIWIS paradigm were surprisingly unsuccessful (Stefik, Bobrow, Foster, Lanning & Tatar 1987). Designers misunderstood the implications of relaxing the WYSIWIS constraints with severe consequences for the users:

“...they found it so frustrating that they put their heads in their hands, raised their voices, and ultimately threatened to walk out. They expressed astonishment that anyone would build such a tool”, Tatar, Foster & Bobrow (1991), page-190.

These difficulties, which manifest themselves in breakdowns such as mistaken deictic reference¹, are largely due to the reduced mutual *awareness* of colleagues' location and activities.

In this paper we investigate the awareness concept in groupware. The design, development, and evaluation of a novel real-time groupware system is used to exemplify awareness issues. Section 2 provides an integrative review of prior work on the models and theories that can be used to analyse awareness requirements and to guide the design of awareness support mechanisms in groupware. The findings of the review are used to motivate our investigation of distortion-oriented mechanisms for supporting collaborative awareness. Sections 3 and 4 discuss the design and implementation of DOME, a Distortion-Oriented Multi-user Editor. We developed DOME to provide an experimentation platform for awareness in groupware, and it demonstrates ways that computer supported information spaces can offer awareness capabilities that have no natural analogy in the real-world. DOME's evaluation (Section 5) reveals fundamental implementation errors that, although obvious in retrospect, provide insights into the over-riding importance of supporting each user's personal focus in the information space, even during collaborative work. Precise formulations for the removal of these errors are presented. Directions for further work are presented in Section 6, and the findings of the paper are summarised in Section 7.

2 Awareness in CSCW

Dourish & Bellotti (1992) defined awareness as “an understanding of the activities of others, which provides a context for your own activity.” This definition captures many user needs ranging from dynamic awareness of colleagues' actions on a shared information space such as a whiteboard through to the more static knowledge about people's status, access privileges to information, prior actions, and so on.

Tang (1991) shows that when working around shared physical work-surfaces such as whiteboards, collaborators move seamlessly between periods of tightly coupled collaborative work with direct awareness of each others activities and periods of more autonomous work. Subtle cues such as gestures and non-verbal utterances are significant signals in the mediation of mutual awareness. People unconsciously monitor these signals and adapt their subsequent actions appropriately.

Distraction is the counter-side to awareness, arising when a person's attention is drawn away from the topic or artifact on which they wish to focus. Computer supported avatars for communicating awareness must do so through constrained media channels which do not readily support subtle and fine-grained awareness cues. They are therefore prone to distracting the users. The effectiveness of awareness support can be seen as a function of the awareness support afforded, over the distraction that it causes.

To maximise the effectiveness of computer supported mechanisms for awareness, it is necessary to understand the underlying factors that influence it. Gutwin (1997) and Benford, Bowers, Fahlen, Mariani & Rodden (1994) provide contrasting models of awareness issues in groupware. Gutwin's model (Section 2.1) attempts to identify the constituent information sources that communicate awareness. His aim is to provide a framework that can assist the evaluation and design of awareness capabilities in groupware. Benford et al.'s model (Section 2.2) is derived from key observations in the social sciences. It focuses on the role of

¹Deictic reference is a combination of reference by gesture and verbal reference such as “this one”.

the information space in mediating awareness. Each of these frameworks is briefly reviewed below.

2.1 Gutwin’s Awareness Framework

Gutwin’s model is derived from a top-down decomposition of awareness types, with particular emphasis on “workspace awareness” which he defines as the awareness of others that is mediated by, or is closely related to, their actions on or around a shared workspace. The model identifies seven types of awareness, but Gutwin argues that four (described below) apply specifically to group work dynamics².

2.1.1 Informal Awareness

Informal awareness concerns the sense of community among a group of colleagues. In real-world office environments this information is gathered by chance causal interactions such as coffee-room discussions and meeting people by the printer. Groupware researchers have investigated ways of removing the pivotal role of close physical proximity in mediating informal awareness. Example systems include “virtual hallways” such as the Cruiser (Root 1988, Fish, Kraut, Root & Rice 1993) and media-spaces such as Portholes (Dourish & Bly 1992), RAVE (Gaver, Moran, MacLean, Lovstrand, Dourish, Carter & Buxton 1992) and the Xerox–Portland link (Bly, Harrison & Irwin 1993). These systems aim to enhance a sense of community over distance by providing video feeds or snapshots of remote offices and communal spaces. Telefreek (Cockburn & Greenberg 1993) investigates ways of providing equivalent awareness information through low-bandwidth text-based media. Active Badge technologies (Want, Hopper, Falcao & Gibbons 1992) can also support informal awareness by automatically gathering information on the location of badge wearers. All these technologies raise many questions of social acceptability (Harper, Lamming & Newman 1992).

2.1.2 Social Awareness

Gutwin describes social awareness as the backchannels of communication that contextualise interaction. In normal face to face conversation, for instance, backchannels include facial expressions, nods and utterances such as “uh-huh”. Backchannels are strongly dependent on the media of interaction, and the paucity of backchannels in textual media helps to explain phenomena such as email flaming.

Several systems have been developed to experiment with enriching the backchannels available in different media, particularly in video streams. Smith, O’Shea, O’Malley, Scanlon & Taylor (1991) and Buxton (1993), for example, explore the use of video-tunnels to support eye-to-eye contact. Ishii, Kobayashi & Grudin (1992) extend this work in their Clearboard system which integrates video-mediated eye-to-eye contact and “gaze awareness” *within* a shared workspace.

2.1.3 Group-structural awareness

Group-structural awareness concerns the protocols and structures that are used to formalise collaboration. Issues such as roles and responsibilities, rights of access to information sources,

²The remaining three awareness types concern the individual’s knowledge about their task and environment.

decision-making processes, and turn-taking in formal meetings are all classified under group-structural awareness.

Many groupware systems have supported a variety of types of group structural awareness. Posner & Baecker (1993) provide a taxonomy of collaborative writing with respect to a variety of issues including the co-writers roles, writing strategies and document control methods. Several systems demonstrate support for group decision making and for the maintenance of organisational memory. These include the Coordinator (Winograd 1987), gIBIS (Conklin & Begeman 1988) and a wide range of workflow support tools such as InConcert (Abbot & Sarin 1994).

2.1.4 Workspace Awareness

Gutwin defines workspace awareness in terms of the capabilities of the media of collaboration: “the awareness mediated by the workspace”. Although this definition allows clear classification of actions on physical work surfaces (such as shared whiteboards), it is more difficult to apply to groupware where the workspace rendered on the screen may provide the only media for collaboration. The groupware workspace must therefore act as a surrogate for all of the other types of awareness described above. The Clearboard system (Ishii & Kobayashi 1992), for example, is designed to integrate social awareness from “gaze-awareness” and “eye-to-eye contact” within the workspace.

More usefully, Gutwin decomposes workspace awareness into a set of constituent information elements which attempt to cover all awareness requirements, including those of the other awareness categories. This articulation of awareness elements, summarised in Table 1, is useful because it supports designers and evaluators in assessing the awareness capabilities of their systems. It also promotes consideration of new ways that groupware could allow users to tailor their awareness of colleagues. For instance, a virtual environment might allow users to attenuate their awareness of colleagues’ activities while accentuating awareness of their location.

Element	Relevant Questions
Identity	Who is participating in the activity?
Location	Where are they?
Activity Level	Are they active in the workspace? How fast are they working?
Actions	What are they doing? What are their current activities and tasks?
Intentions	What are they going to do? Where are they going to be?
Changes	What changes are they making? Where are changes being made?
Objects	What objects are they using?
Extents	What can they see?
Abilities	What can they do?
Sphere of Influence	Where can they have effects?
Expectations	What do they need me to do next?

Table 1: Elements of workspace awareness (extracted from Gutwin *et al.* (1996)).

Greenberg, Gutwin & Roseman (1996) and Gutwin & Greenberg (1998a) describe a variety of awareness “widgets” which use different elements in Table 1 to provide awareness. The

radar view widget, for example, provides a miniaturized rendering of the overall workspace with each user’s location of activity superimposed. A qualitative evaluation of these widgets (Gutwin, Roseman & Greenberg 1996) provided few observations of their effectiveness (or otherwise), despite the subjects’ reported enthusiasm. A subsequent study (Gutwin & Greenberg 1998b) provides preliminary statistical evidence that workspace awareness can improve usability.

2.2 Benford *et al.*’s Spatial Model

Benford et al. (1994) derive their awareness model from an analysis of theories in the social sciences. Their “spatial model of interaction” focuses on the role of the interaction space in mediating collaboration, and it provides a set of abstractions that can be used to characterise and analyse collaborative activities. The key awareness abstractions within the spatial model are “focus” and “nimbus”:

Focus — the more an object is within your focus, the more aware you are of it.

Nimbus — the more an object is within your nimbus, the more aware it is of you.

Figure 1, extracted from (Benford et al. 1994), depicts common real-world awareness modes in collaboration. Mode 5 (“ignoring”), for example, depicts person A projecting their nimbus to encompass B, while B’s focus remains away from A.

Although “focus” is a familiar concept, the term is useful in CSCW because it allows concise description of awareness needs and uses: for instance, a system might allow users to “lock their focus” onto an object, regardless of its location.

The definition of nimbus needs clarification. Nimbus is projected by people and objects. A person who wishes to contribute to the discussion at a meeting, for example, is likely to project their nimbus in order to make others aware of their desire to speak. They may do so by clearing their throat, leaning forward in their seat, waving their hand in the air, and so on. The person intentionally projects their nimbus as an attempt to gain the focus of others. Nimbus may also be projected unintentionally. A whispered side conversation at a meeting, for example, may distract other meeting participants if the talkers are unable to sufficiently suppress their nimbus (the audible murmur of their conversation). The other meeting attendants are unable to satisfactorily filter out the unwanted nimbus.

Nimbus, then, is prone to causing three types of awareness problems.

1. It may be projected unintentionally to people other than those for whom it is intended. This problem arises from the inability of people to adequately tune the ways in which their nimbus traverses a medium (for instance, directing speech at one person to the exclusion of all others) and from the inability of recipients to filter unwanted nimbus.

Groupware has the potential to ease these problems by providing users with tools to accurately target their nimbus.

2. The recipients of nimbus may choose to ignore it, even when it is specifically targeted at them.

This is a social issue for which groupware offers no resolution.

3. The medium may not carry the projected nimbus with sufficient force to gain the focus of the recipient(s). Benford et al. use the term “aura” to describe the range of a




Figure 1: Modes of mutual awareness. Note that the large circles depict the nimbus projected by an object, and the arrows depict the direction of the object’s focus. With permission OUP from Benford et al. (1994).

medium that enables interaction. For instance, consider the differences between the “aura” of the audible medium in a quiet cathedral versus a crowded office environment. Several groupware research projects are investigating ways of reducing (or removing) the impact of remote physical proximity as a factor in collaboration (Section 2.1.1). Essentially these projects seek to create new media with extremely wide aura.

Benford et al. (1994) demonstrate implications of their spatial model within a virtual-reality “populated information terrain”, Q-PIT. In more recent work, Hindmarsh, Fraser, Heath, Benford & Greenhalgh (1998) demonstrate that virtual environments are prone to errors of deictic reference that are similar to those observed in the use of Colab (Tatar et al. 1991).

3 Design Goals for DOME

DOME is a distortion-oriented multi-user text editor that provides a platform for experimentation with workspace awareness. Its development has been guided by four fundamental design goals.

First, we want our system to be able to support realistic collaborative work. We strongly believe that point systems demonstrating technical feasibility, such as those of our previous work (Greenberg, Gutwin & Cockburn 1996a, Greenberg, Gutwin & Cockburn 1996b, Greenberg, Gutwin & Cockburn 1996c), are insufficient for effective evaluation of awareness support. We also believe that small awareness mock-up systems such as the welding simulator used by Gutwin & Greenberg (1998b), although ideal for controlled quantitative experimentation, will provide different insights into awareness needs than those revealed in rich and realistic work support systems. Collaborative text-editing is a natural domain for our system to support because of the range of tasks and collaborative styles that can be investigated.

Second, we assume that single-user acceptability is a pre-requisite to groupware acceptability. If users are unwilling to accept a system for their own use, we see little reason for them accepting it for groupwork (Cockburn & Jones 1995). DOME’s groupware capabilities are seen as enabling collaboration whenever required, but these capabilities do not inhibit the

utility of the system in single-user use.

Third, we want DOME to provide “open” support for collaboration. Rather than support a particular mode of collaboration we want DOME to allow a wide range of collaborative activity, and we want its awareness capabilities to be a major lever in supporting timely and appropriate transitions between collaborative modes. The grain of information access in DOME is the individual text character, and users are free to simultaneously edit the same word if they choose to do so. This open strategy for collaboration has implications that will invalidate DOME’s use in certain collaborative tasks, such as those where issues of information ownership and access privileges are important. We assume that DOME users are trusting collaborators, and we leave it to them to develop their own social protocols for work management. For further discussion on access privileges, granularity and privacy, see Posner & Baecker (1993) and Bellotti (1996).

Finally, we want DOME to allow us to investigate the effectiveness and usability of distortion-oriented mechanisms for supporting collaborative awareness. The reasons for our interest in distortion-oriented awareness schemes are discussed in the following section.

3.1 Distortion-Oriented Awareness Mechanisms

Distortion-oriented visualisation schemes such as magnification lenses, zooming functions and fish-eye lenses (Leung & Apperley 1994) are useful in a wide range of single-user computing systems. They assist visualisation of large information spaces by easing the transition between high-levels of detail in a local area of interest and the global context of the information-space.

Distortion-oriented mechanisms appear to offer additional benefits in groupware environments. They have the potential to allow collaborators to maintain a focussed area of interest in their own work while simultaneously displaying the context of the overall information space. Awareness information such as the location and activities of colleagues can then be superimposed on the work-surface. Furthermore, Dourish & Bellotti (1992) stress that awareness information should be presented in the same information space as the objects of collaboration. Magnification techniques such as fisheye lenses provide a contiguous display of focus and context, and are therefore natural candidates for displaying the individual’s focus and awareness information within a single integrated information space.

We have designed and built DOME to support tailorable distortion-oriented awareness capabilities which we can adapt to create varying experimental conditions and which end users can easily customise to match their awareness needs.

Finally, it is worth noting that although distortion-oriented mechanisms are usually applied within graphical information spaces such as maps and networks, there has been prior research with their use in text-based displays. The paper in which Furnas (1986) first introduced the term “fisheye view” describes a text-based programming environment. Keahey & Marley (1996) reports the evaluation of a single-user distortion-oriented text viewer which displays the entire document within a single window (by reducing the vertical spacing between lines). The user controls a scrollable lens which reveals a legible portion of text with normal vertical spacing. His evaluation indicates that the technique can assist certain searching tasks.

4 DOME: An Awareness Testbed

As a single user system, DOME (Figure 2) is a functionally rich text editor. In addition to normal text-editing utilities such as cut, copy and paste, it includes extensive annotation

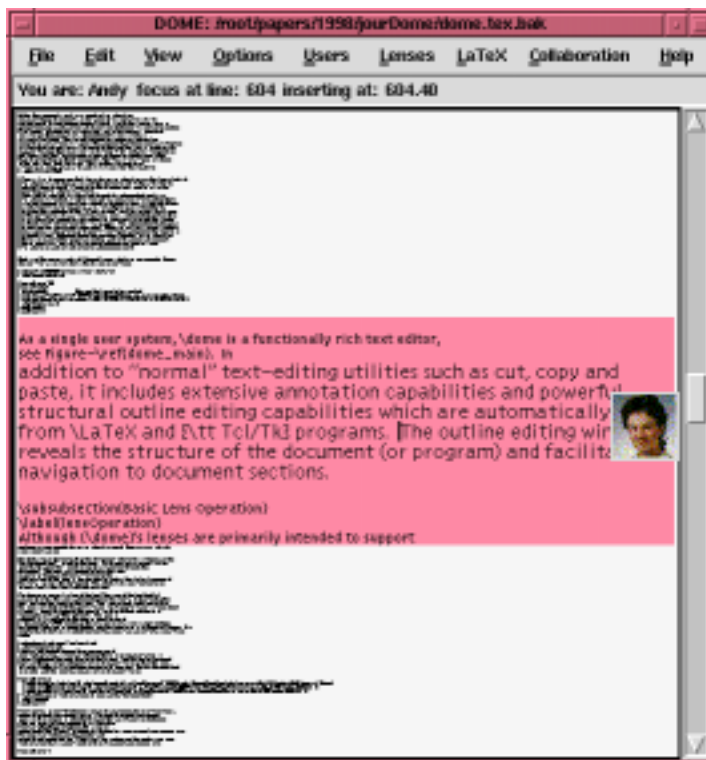


Figure 2: DOME’s main window, showing the focal lens of a single user.

support and powerful structural outline editing capabilities which are automatically parsed from \LaTeX and Tcl/Tk programs. The outline editing window reveals the structure of the document (or program) and facilitates rapid navigation to document sections.

The default size for text within the document, referred to in DOME as the “background text”, is configurable. Small sizes of background text allow large segments of the document to be seen within the display, providing a wide context of the user’s location in the document (see Figure 2). Each user has a magnification lens that follows their insertion point as it moves through the document. The user can choose to have the lens magnify (or demagnify) text according to the configuration of the lens. This allows the user to control the distinctiveness of their focal-point in the document.

When used for simultaneous collaboration, each user has one additional focal point for each of the other users. Awareness of a particular colleague is tailored by configuring the lens associated with that person—from no awareness (no magnification) to high levels of awareness (extensive magnification).

Simultaneous users are assumed to have a continually open audio channel. Each users’ information space is immediately updated to reflect the actions of their partners. Users are responsible for saving their own local copies of the document.

4.1 Basic Lens Operation

Figure 3 shows three examples of lens configuration: in Figure 3(a) there is no magnification; in Figure 3(b) the user has selected a small font for the background text, and there are two



(a) No distortion.

(b) Stepped distortion.

(c) Barrel distortion.

Figure 3: Example configurations of lens and background text.

stepped levels of magnification within the lens; in Figure 3(c), an extreme example of lens use, the background text is minute, and the lens has many levels of magnification.

The degree of magnification of the text (the size of the text font) at any point within a lens is dependent on the vertical distance that the text lies from the centre of the lens. The lens is automatically centred on the user’s insertion point. The text in the lens is tinted with the associated user’s colour, and colours are consistent across all users to remove the possibility of mistaken reference in comments such as “the red text”.

A menu of named awareness levels allows users to rapidly change their awareness of collaborators. Figure 4(a) shows Philip using the menu under the “Users” window to change to a new named awareness level. The number of awareness levels and the names associated with the levels are tailorable, as is the underlying lens configuration associated with each name.

To customise one of the named lens configurations, the user displays the normally hidden “lens handles”. The handles allow three properties of the lens to be modified.

1. *Number of lens steps* — the lens can contain any number of discrete magnification steps. In Figure 3(a) there are zero magnification steps; there are two steps in Figure 3(b); and there are six steps in Figure 3(c).
2. *Magnification at each step* — the text size within each step is set by dragging lens handles horizontally (rightwards for larger, leftwards for smaller).
3. *Range of each step* — the number of lines of text within each step is controlled by dragging the lens handles vertically.

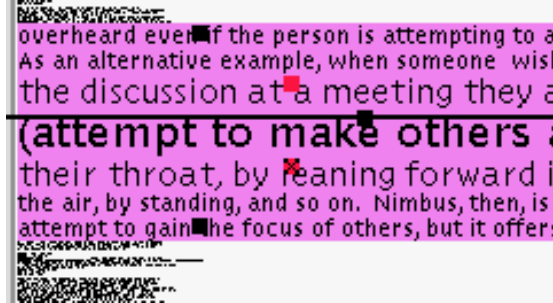
Figure 4(b) shows the user customising a lens with three magnification steps: the thick horizontal line shows the centre of the lens, and the dots are the lens handles.

4.2 Fisheyes for Tailorable Workspace Awareness

Figure 5 shows two users’ simultaneous views of a document. Figure 5(a) shows that Andy has dedicated approximately a third of the display space to his own focal region (at the top



(a) Awareness levels menu.



(b) Customising lens configuration with lens handles.

Figure 4: Tailoring Awareness.

of the display). He has a roughly paragraph size lens onto Philip’s work at a just-legible text size, and the background text is set to a very small (not legible) size which gives an extensive overall view of the “shape” of the document. In Figure 5(b), Philip has selected a just-legible text size for the background text, and a slightly larger text-size for his own focal point. His lens onto Andy’s work is configured to show location only, with no magnification of the background text. Unlike Andy, he has chosen not to view the user’s image at the centre of each lens.

The size of the background text is an important determinant in the maintenance of collaborative awareness. In a large document, even with a small size of background text, a user who’s focus is towards the end of the document will not be able to see the actions of a user who’s focus is at the start of the document because their region will be scrolled out of the display. Benford et al. (1994) used the term “aura” to describe the range of the medium that allowed awareness to be communicated (Section 2.2). The smaller the background font-size the greater the aura of the medium. DOME allows the background text to be set to a zero (invisible) font-size. Users can then guarantee that collaborators’ focal points do not scroll out of the visible text region by configuring the lenses so that their total height is the same as that of the window: the entire document is then within the medium’s aura.

4.3 Nimbus controls and customising awareness

DOME’s lenses allow users to control the nimbus that others project. This is unlike the physical world where control over others’ nimbus is hard to achieve. Although this may be advantageous in some collaborative activities, it is also likely to frustrate users who urgently need to gain the attention of others.

DOME provides several interface mechanisms that support the projection of nimbus. First, each user’s cursor location is constantly transmitted to all other users through semantic tele-



(a) Andy's view.



(b) Philip's view.

Figure 5: Two simultaneous views of the workspace.

pointers which are represented as large dots over the text. The coordinates of the telepointers are adjusted to ensure that they lie over the same item of text for each user. This helps to assure that deictic references such as “This sentence” (with a gesture around the associated text) can be correctly interpreted regardless of the different window sizes that users may have. Each user’s text-selection is also transmitted to all other users, and is represented by shading the text in that user’s colour. Neither of these nimbus mechanisms will succeed if the associated text region is scrolled out of the colleague’s window.

Second, users can move their focal point to the location of another user by selecting “Goto” from the user’s menu. This method offers no guarantee of successfully projecting nimbus because the targetted user can select a zero awareness level for specific users.

The third mechanism is an intrusive one that is intended to help users who want to initiate tightly coupled interaction. Any user can issue a “Follow” request to another user. This causes a dialogue-box to appear on the recipient’s screen, who must then accept or decline the request. If accepted, the requester’s window is re-configured to assure strict-WYSIWIS between the two windows, and all scrolling actions by the “followed” user are immediately replicated in the follower’s window.

Annotations, described in the following section, provide the final mechanism for supporting focus and nimbus.

4.4 Annotations

DOME’s annotations follow a sticky-note metaphor. Annotations can be associated with a particular piece of text (“stuck onto the paper”), or they can float above the page (“stuck to the screen”) with a constant geometry while the page is scrolled “underneath”. They may



Figure 6: Three annotations.

also be either private or shared, with shared annotations being simultaneously editable by all users. Annotations that are initially created as private can later be shared, and they can be sent to specific people or subgroups.

DOME’s annotation capabilities are intended to satisfy several user needs.

1. *Temporal awareness and asynchronous collaboration.* Because annotations persist between system invocations, they can be used to provide temporal awareness of the changes made during periods of asynchronous work.
2. *Marginal notes and meta-comments.* It is common for authors to make marginal notes. Private notes can be used as personal reminders about pending actions on the document, and shared notes can be used to discuss topics associated with specific pieces of text.
3. *Dynamic Awareness and communication portholes.* Shared annotations can be used to provide continually active communication portholes.

Figure 6 shows three annotations. In the top-right corner there is a floating shared annotation serving as a general communication porthole. In the bottom left corner there is a shared annotation that is pinned to the text, and in the middle of the screen is a private-pinned annotation that the user is about to send to user “Andy”.

5 Evaluation

Our motivation for building DOME was to evaluate awareness issues, particularly the utility of distortion-oriented mechanisms for awareness, within realistic work tasks. From our personal experiences in using DOME³, we felt we had constructed a well-considered and polished system that could reasonably be used to support collaborative authoring.

A heuristic evaluation of the potential utility of the system with respect to Gutwin’s constituent information elements of awareness (Table 1) is shown in Table 2.

³An earlier version of this paper was written using DOME (Weir & Cockburn 1998).

Element	Supported & How	Comments & Problems
Identity	<ul style="list-style-type: none"> ✓ User image at lens-center. ✓ Colour of focal region. 	<ul style="list-style-type: none"> ✗ Image at centre of lens may be turned off. ✗ 1. Focal point may be scrolled out of view. ✗ 2. Lens may be configured to give no awareness. ✗ Possible problems mapping from colour to identify
Location	<ul style="list-style-type: none"> ✓ Location of lens. ✓ Semantic telepointers. ✓ Strict-WYSIWIS through “Follow” utility. 	<ul style="list-style-type: none"> ✗ See 1&2 above. ✓ Zero size for background text assures lens location within display, but reduces context. ✗ See 1&2 above. ✗ Strict-WYSIWIS may be too constraining.
Activity Level	<ul style="list-style-type: none"> ✓ Actions visible through lens or background. 	<ul style="list-style-type: none"> ✗ See 1&2 above.
Actions	<ul style="list-style-type: none"> ✓ Actions visible through lens or background. 	<ul style="list-style-type: none"> ✗ Low awareness levels or small background may not reveal details. ✗ See 1&2 above.
Intentions	<ul style="list-style-type: none"> ✓ Lens movement may indicate intentions. 	<ul style="list-style-type: none"> ✗ See 1&2 above.
Changes	<ul style="list-style-type: none"> ✓ Dynamic updates visible in lens or background. ✓ Annotation facilities. 	<ul style="list-style-type: none"> ✗ See 1&2 above. ✗ Low awareness levels and small background text may not reveal changes. ✓ Annotations persist for asynchronous collaboration updates. ✗ Only available through explicit user action.
Objects	N/A	The document is the only object.
Extents	<ul style="list-style-type: none"> ✓ Strict-WYSIWIS forces shared views. 	<ul style="list-style-type: none"> ✗ Strict-WYSIWIS is highly constraining.
Abilities	N/A	No granularity or access control.
Sphere of Influence	<ul style="list-style-type: none"> ✓ As for Location above. 	
Expectations	<ul style="list-style-type: none"> ✓ Annotations 	Explicit meta-actions can be recorded in annotations.

Table 2: Assessment of DOME with respect to Gutwin’s workspace awareness elements (Table 1).



Figure 7: Co-located seating arrangements in the evaluation.

5.1 Evaluation Method

The evaluation was intended to detect fundamental usability flaws in the system, and to make indicative observations on the utility of the distortion-oriented awareness mechanisms. Prior groupware research has repeatedly reported systems that fail to measure up to the designers' expectations (Grudin 1990), so we believed it prudent to carry out a broad heuristic usability evaluation of DOME before initiating evaluation based on quantitative measures of process and task performance.

The fourteen subjects in the evaluation were unpaid postgraduate Computer Science students (12) and Computer Science academics (2). All were familiar with the concepts of synchronous groupware, but none had previous experience with DOME. Given the goal of the evaluation—to detect usability problems and make indicative observations—we were not concerned about the obvious subject-bias arising from the use of groupware familiar Computer Scientists. The subjects were randomly assigned to pairs. Each pair participated in a single one-hour video-taped session. The evaluators were seated side-by-side approximately one and a half meters apart. Figure 7 shows the seating positions in the evaluation. The co-located side-by-side seating position was selected for two reasons. First, it illuminated failings in DOME's awareness support: if collaborators found it necessary to look at their partner's screen, it would be a strong indication that the system's awareness facilities were insufficient. Second, it greatly eased the evaluation process: a single evaluator and a single video-stream could observe and capture all user actions. The side-by-side arrangement reduced the realism of the evaluation, but it did not detract from the coarse-grained usability objectives of the study.

The first twenty five minutes of each session were spent introducing DOME's capabilities, with emphasis on the operation of the fisheye lenses. Although this is a long time relative to the total length of the evaluation session, we felt that it was important to attempt to thoroughly familiarise the users with the lens operation prior to making observations on use. We emphasised that awareness of partners was best controlled through the menu of named awareness states.

The pre-configured system state for all users was the same, with seven named awareness

levels, ranging from “none” (no distortion or representation of the lens) through to “huge” (24 lines of large text). At the start of the session the background text was set to a just-legible size (6 point), the partner’s nimbus provided a legible region of 6 lines at an 8 point font, and the user’s own lens provided 10 lines at a 10 point font.

Fifteen minutes were then spent on each of two tasks, and five minutes were allocated to a post task interview. The first task involved re-ordering a jumbled structured document. It was introduced to the users as follows:

A production of “The Tragedy of Macbeth” by William Shakespeare is due to be performed tonight. The problem is that the script was emailed to the production company in parts, and assembled automatically by an email reassembler. The software malfunctioned and assembled the messages based on the order in which they arrived, rather than the sequence in which they were sent.

Your task is to work together with your partner to reassemble the text so the play can go on. The file `macbeth.tex` contains the jumbled text.

In the second task the subjects were given a template document and were asked to complete the sections, and in particular to summarise these sections in a table at the end of the template. The task was introduced as follows:

Your second task is to fill-in the details of a document on the fourth year Computer Science papers offered in our department. The structure of the paper is provided in the file `course.tex`. Be sure to complete the summary table at the end of the document.

The two tasks were selected to generate differing requirements for workspace awareness. The first task was intended to promote frequent transitions between collaborative work and independent work as the users moved from negotiating which scene to collect next (which we presumed would be a collaborative activity) to retrieving and repositioning the text (which we presumed would be an independent activity). The second task was selected to generate a need for workspace awareness in more tightly coupled collaborative activity. Although some of the sections in the template could be completed independently, we believed it would be necessary to collaborate in order to produce the summary table.

5.2 Observations

Although all of the pairs managed to successfully complete the tasks, the overall reaction to the distortion-oriented awareness features was negative. The subjects reported, and we observed, some substantial limitations in the implementation of the fisheye lenses. These problems, which dominated the evaluation, are described in Sections 5.3 and 5.4 together with methods for correcting them.

Not all of the observations were negative, and there were many observed events indicating that the users were successfully maintaining awareness of one another. For instance, during the second task which involved writing about courses and lecturers, we noted one of the subjects laughing. When asked, the subject reported that he had seen what his partner was writing about a lecturer. Although the information communicated was frivolous (in this case), it clearly indicates that the user was aware of their partner’s activity. Several of the pairs also made successful use of the shared annotations for mediating collaborative activity, such

as keeping track of assigned tasks. Despite the side-by-side location, subjects seldom looked at their partner’s screen.

One pair was particularly enthusiastic about the awareness capabilities of the system. Early in the session they configured the background text to an invisible font-size, and selected large awareness regions for their own focus and for their partner’s nimbus. This configuration assured that their partner’s focus would always be visible, regardless of their location in the document. Both users reciprocated in their lens settings, assuring that they had equivalent awareness levels of each other. They made no subsequent modifications to their awareness settings. Like all other pairs they stated that they were disturbed by the jumping text (Section 5.3), but they commented that “it was good to be able to see both areas of interest”.

Only two other pairs made any alteration to the default awareness settings. One of these turned off all awareness by selecting the “none” awareness level. Both users simultaneously set their systems to this configuration. They reported that the absence of awareness information allowed them to “focus on the task at hand”. Their collaboration was characterised by very little speech, with both users essentially attempting to solve the tasks independently. Periodically they would change their independent sub-tasks on noticing that their colleague had already accomplished the task that they intended to do next. The other pair that modified the lens configurations made several changes through the session. They were the only pair to use non-reciprocated lens settings. Their collaboration was also characterised by little speech and mainly individual work.

5.3 “Jumping” text problems

The biggest problem, reported by all of the subjects, was caused by “jumping text”. In specific circumstances (described below) the position of text at the user’s focal point would move by up to 3cm (in the worst observed case), requiring the user to find the new location of their insertion point. Similar problems of absent ‘display inertia’ have been observed in research into holoprasting systems (Barstow 1984, Smith, Barnard & Macleod 1984) that allow users to tailor the contents of textual displays. In DOME the problem is exacerbated because text-motion can be caused by any of the simultaneous users of the system. Although the errors were clear once evaluation was underway, it is interesting that they had not been detected previously: the same errors existed in an earlier prototype system (Greenberg, Gutwin & Cockburn 1996*b*) which had been extensively demonstrated by the designers, but *not* evaluated by independent usability subjects.

The text jumps occur under two conditions. First, when a user relocates their insertion point (or I-beam) the text around that point is magnified, causing an offset between the x , y coordinates of the click and the resultant position of the I-beam. The size of the offset depends on the difference between the initial size of the text (normally the background text) and the final magnified size of the text. Equations 1 and 2 precisely express the magnitude and direction of the offset. Figure 8 demonstrates the problem: the cross-hair in the left-hand text-segment shows the coordinates of the user’s click, and right-hand text-segment shows the offset to the I-beam. Although several users indicated surprise on first seeing this effect, they quickly became used to looking for the nearby location of the I-beam having repositioned their lens. The blinking of the I-beam also helped draw the user’s eye to its new location.

Second, whenever a colleagues’ lens enters or leaves the visible text region above the local user’s lens, the change in text size causes a vertical offset in the local user’s view. The effect is most noticeable at high levels of awareness, and when the background text is set to a small

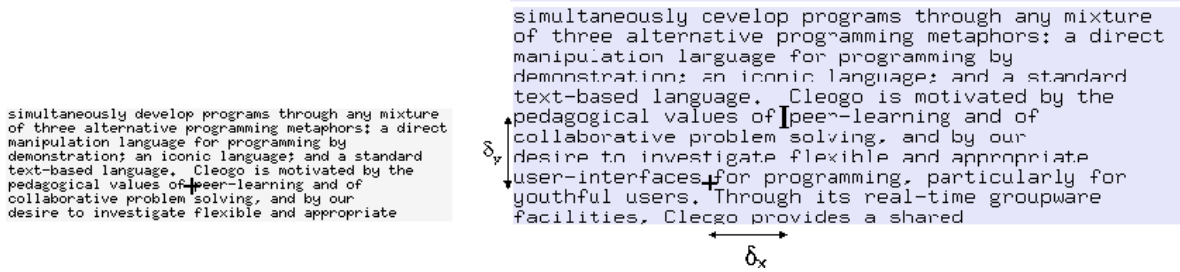


Figure 8: Unwanted text motion when repositioning the insertion point. The clicked position is shown on the left, and the offset to the I-beam is shown on the right.

size. Equation 3 shows a variant of Equation 2 that accounts for and corrects the δ_{y0} offset caused by lens relocation. This cause of text-motion was more disturbing for users because it arose “unpredictably” as a consequence of the partner’s movement within the document.

The precise details of the occurrence of this problem are partially dependent on the behaviour of the graphical user interface toolkit and widget set used to implement the lens mechanisms. However, the magnitude of the offset as shown by Equation 3 is independent of toolkit specifics.

$$\delta_x = \sum_{c=0}^{clickChar} w_R(c) - w_I(c) \quad (1)$$

where

- δ_x is the horizontal offset between the clicked location and the resultant location of the I-beam.
- c takes the value of each character from the start of the line to the clicked character.
- $w_R(x)$ gives the width of character x at the resultant (magnified) font-size. Note that this value is the same for all characters in fixed-width fonts.
- $w_I(x)$ gives the width of character x at the initial (background) font-size.

Equation 1. Horizontal offset between the clicked location and the resultant location of the I-beam.

$$\delta_y = \pm \left((h(1) - h(b)) * \frac{lines(1)}{2} + \sum_{l=2}^{numLevels} (h(l) - h(b)) * lines(l) \right) \quad (2)$$

where

- δ_y is the vertical offset between the clicked location and the resultant location of the I-beam. δ_y is negative (the I-beam appears below the clicked location) when the user clicks a text location above the current location of the fisheye. It is positive when the user clicks below the fisheye.
- $h(x)$ gives the font-height at magnification level x .
- b is the background magnification level.
- l takes one value for each of the stepped lens levels of magnification, from the 2nd stepped level to the outer-most level. Level 1 is the lenses centre of magnification.
- $lines(x)$ gives the number of lines contained within lens magnification level x .

Equation 2. Vertical offset between the clicked location and the resultant location of the

I-beam.

$$\delta_{yo} = \pm \left((h(1) - h(b)) * lines(1) + \sum_{l=2}^{numLevels} (h(l) - h(b)) * 2 * lines(l) \right) \quad (3)$$

where

- δ_{yo} is the vertical offset caused by other’s motion around a user. δ_{yo} is positive (the text jumps upwards) when the colleague moves from above the user to below. It is positive when the colleague moves from below to above.
- $h(x)$, b , l , and $lines(x)$ are the same as in Equation 2.

Equation 3. Vertical offset caused by a colleague’s movement around a user.

5.4 Location versus Semantic Distortion

We observed several occurrences of subjects becoming disoriented in the document, typified by accompanying comments such as “Now, where’s the sub-section gone?” There were apparent periods of conflict between the text that the subjects expected to see magnified and the text that was magnified. The problem was particularly pronounced when the subjects had selected extremely small background text-sizes, and when section or subsection headings rolled out of the magnified region of the lens.

DOMÉ’s lenses magnify text purely according to the vertical distance that the text lies from the centre of one of the users’ lenses. The distortion functions do not account for the semantic importance of the underlying text. Furnas (1986) used the term “*a-priori* importance” to capture this notion of semantic importance within information spaces. In text documents there is a high level of *a-priori* importance associated with the structural tags of the document: section headings, titles, captions and so on. Users who are familiar with WYSIWIS word processors will expect to see these tags displayed in a format that is similar to their final rendering—normally large and emboldened text.

6 Discussion and Further Work

Although the problems caused by unwanted text motion overwhelmed our evaluation, we are not overly discouraged by our observations. We witnessed many positive uses of the awareness information provided by DOMÉ, and one of the pairs was enthusiastic about the awareness provided by the distortion-oriented lenses despite the limitations of the current implementation.

The most severe usability problems observed were derived from the unwanted δ_{yo} offset. Users clearly did not want the location of their focal point to be affected by colleagues’ actions. These problems could be resolved through several alternative modifications in DOMÉ. The simplest would be to introduce a granularity of updates into DOMÉ’s currently fully synchronous and zero granularity model: for instance, colleagues’ actions could be transmitted to others only when the local user requests them or when the user scrolls to a new region in the document. This option would be a severe departure from our initial objectives for DOMÉ because it would remove much of DOMÉ’s support for dynamic and continuous awareness information.

The remedy that we have selected for our next version of DOME is to use Equations 2 and 3 to correct the δ_{yo} and δ_y offsets. These equations provide pixel displacement values which can be readily converted into values such as line-numbers or coordinates which are used to control scroll locations in most common graphical user interface toolkits. We will not correct horizontal displacement (δ_x offsets). To do so would require either horizontal scrolling, enlarged margins, or modifications to the line-wrap. We doubt that any of these alternatives would be a significant improvement over the relatively small δ_x offsets. Horizontal scrolling would result in part of the text of the line being obscured, with obvious detrimental effects on readability. Enlarged margins would require the width of the window to change without a direct command from the user. Modifications to the line-wrap would disrupt the user's text formatting. Subsequent evaluations will determine whether fixing the vertical offsets is sufficient for user acceptance.

Another problem detected in the evaluation that we will not address in our next version of DOME is the "location versus semantic content" problem. To date our interests have been in collaborative awareness, and the specific details of the media in which we investigate collaborative awareness are secondary. Prior to addressing system refinements that are closely associated to the media (the semantics associated with specific tags in the document), we wish to gain a better understanding of the large grained awareness issues. In the meantime, we will use DOME's structural outline viewer for evaluation tasks that require a semantic representation of the workspace.

Having corrected the errors in our implementation, we will continue to investigate the boundaries between awareness facilitation and distraction. We wish to undertake longitudinal studies to investigate how collaborators use the distortion-oriented awareness mechanisms once they are thoroughly familiar with the system. Such an investigation is clearly dependent on successful results in a preliminary evaluation such as that reported in this paper.

7 Conclusions

Awareness of colleagues' activities has been recognised as a pivotal factor in collaborative work. In this paper we reviewed the awareness models and frameworks that can be used to guide the design of groupware tools that support real-time collaborative awareness.

Drawing on these models and frameworks, it appears that distortion-oriented visualisation techniques have the potential to satisfy collaborators' needs for workspace awareness in real-time groupware environments. They allow users to focus on their local region of interest while simultaneously providing a viewport onto the work of colleagues.

Our distortion-oriented multi-user editor, DOME, has been designed to provide a fully-functional environment for experimenting with and evaluating distortion-oriented workspace awareness. Our intention is to go beyond demonstrational point systems and to experiment with awareness support mechanisms within realistic collaborative work tasks. The first evaluation of DOME illuminated several substantial usability problems with our implementation of awareness capabilities through "fisheye" distortion. The evaluation revealed that users are extremely sensitive to violations in the visual representation of their focus of interest. The fact that several designers failed to detect these flaws in DOME and in its prototype, despite extensive "internal" testing and many public demonstrations, provides further confirmation of the frailty of designers' intuition (and tunnel vision) in designing and evaluating groupware systems (Tatar et al. 1991, Grudin 1990).

Although these problems overwhelmed the evaluation results, we are encouraged by evidence of successful uses of the awareness features. Our further work will ease the usability problems and continue evaluation.

Availability and Acknowledgements

DOMÉ is written in Tcl/Tk (Ousterhout 1993) and GroupKit (Roseman & Greenberg 1996). Its starting point was a fisheye text viewer developed by Saul Greenberg while on sabbatical at the University of Canterbury. DOMÉ has been tested on Sun Sparc stations and on PCs running the Linux operating system. It is available on request from the first author.

References

- Abbot, K. & Sarin, S. (1994), Experiences with workflow management: Issues for the next generation, *in* ‘Proceedings of CSCW’94 ACM Conference on Computer Supported Cooperative Work, October 22–26. Chapel Hill, North Carolina.’, pp. 113–120.
- Barstow, D. (1984), A display-oriented editor to interlisp, *in* D. Barstow, H. Shrobe & E. Sandewall, eds, ‘Interactive Programming Environments’, McGraw-Hill, pp. 288–299.
- Bellotti, V. (1996), What you don’t know can’t hurt you: Privacy in collaborative computing, *in* ‘People and Computers XI (Proceedings of the 1996 British Computer Society Conference on Human-Computer Interaction.) Imperial College, London. 20–23 August’, Springer-Verlag, pp. 241–261.
- Benford, S., Bowers, J., Fahlen, L., Mariani, J. & Rodden, T. (1994), ‘Supporting cooperative work in virtual environments’, *The Computer Journal* **37**(8), 653–668.
- Bly, S., Harrison, S. & Irwin, S. (1993), ‘Media spaces: Video, audio, and computing’, *Communications of the ACM* **36**(1), 28–47.
- Buxton, W. (1993), Telepresence: Integrating shared task and person spaces, *in* R. Baecker, ed., ‘Readings in Groupware and Computer Supported Cooperative Work’, Morgan-Kaufmann, pp. 846–852.
- Cockburn, A. & Greenberg, S. (1993), Making contact: Getting the group communicating with groupware, *in* ‘Proceedings of COOCS ’93. The ACM Conference on Organisational Computing Systems. November 1st to 4th. Milpitas, California.’, pp. 31–41.
- Cockburn, A. & Jones, S. (1995), ‘Four principles for groupware design’, *Interacting with Computers: the interdisciplinary journal of human-computer interaction* **7**(2), 195–210.
- Conklin, J. & Begeman, M. (1988), gIBIS: a hypertext tool for team design deliberation, *in* ‘Hypertext ’87. University of North Carolina, Chapel Hill, North Carolina, November 13–15 1987’, pp. 247–251.
- Dourish, P. & Bellotti, V. (1992), Awareness and coordination in shared workspaces, *in* ‘Proceedings of CSCW’92 ACM Conference on Computer Supported Cooperative Work, October 31–November 4. Toronto, Canada.’, pp. 107–114.

- Dourish, P. & Bly, S. (1992), Portholes: Supporting awareness in a distributed work group, *in* ‘Proceedings of CHI’92 Conference on Human Factors in Computing Systems Monterey, May 3–7’, Addison-Wesley, pp. 541–547.
- Fish, R., Kraut, R., Root, R. & Rice, R. (1993), ‘Video as a technology for informal communication’, *Communications of the ACM* **36**(1), 48–61.
- Furnas, G. (1986), Generalized fisheye views, *in* ‘Human Factors in Computing Systems III. Proceedings of the CHI’86 conference.’, Amsterdam; North Holland/ACM, pp. 16–23.
- Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. & Buxton, W. (1992), Realizing a video environment: Europarc’s RAVE system, *in* ‘Proceedings of CHI’92 Conference on Human Factors in Computing Systems Monterey, May 3–7’, Addison-Wesley, pp. 27–35.
- Greenberg, S., Gutwin, C. & Cockburn, A. (1996*a*), ‘Applying distortion-oriented displays to groupware’, Video Programme of CSCW’96: The ACM Conference on Computer Supported Cooperative Work.
- Greenberg, S., Gutwin, C. & Cockburn, A. (1996*b*), Awareness through fisheye views in relaxed-wysiwi groupware, *in* ‘Proceedings of Graphics Interface Conference. 21–24 May, Toronto, Canada.’, Morgan-Kaufmann, pp. 28–38.
- Greenberg, S., Gutwin, C. & Cockburn, A. (1996*c*), Using distortion-oriented displays to support workspace awareness, *in* ‘People and Computers XI (Proceedings of the 1996 British Computer Society Conference on Human-Computer Interaction.) Imperial College, London. 20–23 August’, Springer-Verlag, pp. 299–314.
- Greenberg, S., Gutwin, C. & Roseman, M. (1996), Semantic telepointers for groupware, *in* ‘OzCHI’96: The Sixth Australian Conference on Computer-Human Interaction. Hamilton, New Zealand. November 24–27.’, IEEE Computer Society Press, pp. 54–61.
- Grudin, J. (1990), Groupware and cooperative work: problems and prospects, *in* B. Laurel, ed., ‘The Art of human-computer interface design’, Addison-Wesley.
- Gutwin, C. (1997), Workspace Awareness in Real-Time Distributed Groupware, PhD thesis, Department of Computer Science, University of Calgary, Canada.
- Gutwin, C. & Greenberg, S. (1998*a*), Design for individuals, design for groups: Tradeoffs between power and workspace awareness, *in* ‘Proceedings of CSCW’98 ACM Conference on Computer Supported Cooperative Work, November 14–18. Seattle, Washington.’, pp. 207–216.
- Gutwin, C. & Greenberg, S. (1998*b*), Effects of awareness support on groupware usability, *in* ‘Proceedings of CHI’98 Conference on Human Factors in Computing Systems Los Angeles, April 18–23’, pp. 511–518.
- Gutwin, C., Roseman, M. & Greenberg, S. (1996), A usability study of awareness widgets in a shared workspace groupware system, *in* M. Ackerman, ed., ‘Proceedings of CSCW’96 ACM Conference on Computer Supported Cooperative Work, November 16–20. Boston, Massachusetts.’, pp. 258–267.

- Harper, R., Lamming, M. & Newman, W. (1992), 'Locating systems at work: Implications of the development of active badge applications', *Interacting with Computers: the Interdisciplinary Journal of Human-Computer Interaction* **4**(3), 343–363.
- Heath, C. & Luff, P. (1992), 'Collaboration and control: Crisis management and multimedia technology in london underground line control rooms', *Computer Supported Cooperative Work* **1**, 64–94.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S. & Greenhalgh, C. (1998), Fragmented interaction: Establishing mutual orientation in virtual environments, in 'Proceedings of CSCW'98 ACM Conference on Computer Supported Cooperative Work, November 14–18. Seattle, Washington.', pp. 217–226.
- Ishii, H. & Kobayashi, M. (1992), Clearboard: A seamless medium for shared drawing and conversation with eye contact, in 'Proceedings of CHI'92 Conference on Human Factors in Computing Systems Monterey, May 3–7', Addison-Wesley, pp. 525–532.
- Ishii, H., Kobayashi, M. & Grudin, J. (1992), Integration of inter-personal space and shared workspace: Clearboard design and experiments, in 'Proceedings of CSCW'92 ACM Conference on Computer Supported Cooperative Work, October 31–November 4. Toronto, Canada.', pp. 33–42.
- Keahey, T. & Marley, J. (1996), Viewing text with non-linear magnification: An experimental study, Technical report, Computer Science, 215 Lindley Hall, Indiana University.
*<http://ftp.cs.indiana.edu/hyplan/tkeahey/research/papers/TR459.html>
- Leung, Y. & Apperley, M. (1994), 'A review and taxonomy of distortion-oriented presentation techniques', *ACM Transactions on Computer Human Interaction* **1**(2), 126–160.
- Ousterhout, J. (1993), *An Introduction to Tcl and Tk*, Reading, MA: Addison-Wesley.
- Posner, I. & Baecker, R. (1993), How people write together, in R. Baecker, ed., 'Groupware and Computer-Supported Cooperative Work', Morgan-Kaufmann, pp. 239–250.
- Root, R. (1988), Design of a multi-media vehicle for social browsing, in 'Proceedings of CSCW'88 Second Conference on Computer Supported Cooperative Work, Portland, Oregon, September 26–28', pp. 25–38.
- Roseman, M. & Greenberg, S. (1996), 'Building real time groupware with GroupKit, a groupware toolkit', *ACM Transactions on Computer-Human Interaction* **3**(1), 66–106.
- Smith, R., O'Shea, T., O'Malley, C., Scanlon, E. & Taylor, J. (1991), Preliminary experiments with a distributed, multi-media, problem solving environment., in J. Bowers & S. Benford, eds, 'Studies in Computer Supported Cooperative Work: Theory, Practice and Design', North-Holland, pp. 31–47.
- Smith, S., Barnard, D. & Macleod, I. (1984), 'Holophrasted displays in an interactive environment', *International Journal of Man-Machine Studies* **20**, 343–355.
- Stefik, M., Bobrow, D., Foster, G., Lanning, S. & Tatar, D. (1987), 'WYSIWIS revised: Early experiences with multiuser interfaces', *ACM Transactions on Office Information Systems* **5**(2), 147–167.

- Tang, J. (1991), 'Findings from observational studies of collaborative work', *International Journal of Man-Machine Studies* **34**, 143–160.
- Tatar, D., Foster, G. & Bobrow, D. (1991), 'Design for conversation: Lessons from cognoter', *International Journal of Man-Machine Studies* **34**(2), 185–209.
- Want, R., Hopper, A., Falcao, V. & Gibbons, J. (1992), 'The active badge location system', *ACM Transactions on Information Systems* **10**(1), 91–102.
- Weir, P. & Cockburn, A. (1998), Distortion-oriented workspace awareness in DOME, in 'People and Computers XII (Proceedings of the 1998 British Computer Society Conference on Human-Computer Interaction.) Sheffield Hallam University, Sheffield.', Springer-Verlag, pp. 239–252.
- Winograd, T. (1987), 'A language/action perspective on the design of cooperative work', *Human-Computer Interaction* **3**(1), 3–30.