Gutwin, C., and Greenberg, S. (2004) The Importance of Awareness for Team Cognition in Distributed Collaboration. In E. Salas and S. M. Fiore (Editors) *Team Cognition: Understanding the Factors that Drive Process and Performance*, pp. 177-201, Washington:APA Press.

The Importance of Awareness for Team Cognition in Distributed Collaboration

Carl Gutwin[1] and Saul Greenberg[2]

[1]University of Saskatchewan, Saskatoon, Saskatchewan, S7N 5A9, Canada
[2]University of Calgary, Calgary, Alberta, T2N 1N4, Canada
carl.gutwin@usask.ca, saul@cpsc.ucalgary.ca,

Introduction

Although the phrase *team cognition* suggests something that happens inside people's heads, teams are very much situated in the real world, and there are a number of activities that have to happen out in that world for teams to be able to think and work together. This is not just spoken communication. Depending on the circumstances, effective team cognition includes activities such as using environmental cues to establish a common ground of understanding, seeing who is around and what they are doing, monitoring the state of artifacts in a shared work setting, noticing other people's gestures and what they are referring to, and so on (Clark, 1996; Hutchins, 1996).

In this chapter, we argue that *awareness* of other group members is a critical building block in the construct of team cognition, and consequently that *computational support for awareness* in groupware systems is crucial for supporting team cognition in distributed groups. Our main message is that:

> … for people to sustain effective team cognition when working over a shared visual workspace, our groupware systems must give team members a sense of *workspace awareness*.

We describe the collaborative situations we address in this chapter, and then we introduce workspace awareness and discuss why it is a problem in conventional groupware systems.

Collaborative situations

In this chapter, we consider only a subset of collaborative situations. These constrain collaboration to the environment that people work within, the type of systems they use to support distributed collaboration, the tasks that people do, and the type of groups.

- ***Environment: shared workspaces.*** Many teams often work over a shared visual workspace: a bounded space where people can see, generate, and manipulate artifacts related to their activities. We concentrate on flat, medium-sized surfaces (e.g., a large table) upon which objects can be placed and manipulated, and around which a small group of people can collaborate.
- ***Systems: real-time distributed groupware.*** Real-time distributed groupware systems allow teams to work together at the same time, but from different places (e.g. Ellis, Gibbs, & Rein,1991). Here, we are interested only in groupware that provide an electronic equivalent of shared workspace.
- ***Tasks: generation and execution.*** Primary task types in shared workspaces are generation and execution activities (McGrath 1984) in which people create new artifacts, navigate through a space of objects, or manipulate existing artifacts.
- ***Groups: small groups and mixed-focus collaboration.*** Small groups of between two and five people primarily carry out tasks in these medium-sized workspaces. These groups often engage in mixed-focus collaboration, in which people shift frequently between individual and shared activities during a work session (e.g. Dourish & Bellotti, 1992; Salvador, Scholtz & Larson,, 1995).

Within these boundaries, a rich variety of small-group collaboration is possible. Typical real life examples might include two people arranging, ordering, and sorting slides on a light table; a research group generating ideas on a whiteboard; managers of a project planning a task timeline; or a group laying out a page for typesetting.

Workspace awareness and the failings of groupware

Team cognition happens fairly naturally when people work face-to-face over shared physical workspaces. While we recognize that certain task domains may require people to follow an explicit process, people's actions as they perform rudimentary workspace operations are typically graceful and are executed with little conscious effort. An effective team maintains and updates a shared mental model (e.g., Cannon-Bowers, Salas & Converse, 1993) of its members' actions and task work; this happens naturally as each member tracks the evolution of the product developed within the workspace.

All this works so well in face-to-face settings because people easily maintain a sense of *workspace awareness*. We define workspace awareness as:

> the up-to-the-moment understanding of another person's interaction with the shared workspace (Gutwin & Greenberg, 2002).

We elaborate on this definition in later sections, but for now we say that workspace awareness is limited to those activities happening within the temporal and physical bounds of the task that the group is carrying out over a visual workspace. This includes awareness of people, how they interact with the workspace, and the events happening within the workspace.

Workspace awareness is something people take for granted in the everyday world. Because acquiring awareness information is so simple, people rarely consider it as an intentional activity. As a consequence, the role of awareness is often overlooked when analyzing team behavior. In turn, this has meant that groupware systems developed for distributed teams working over some type of shared visual surface—electronic whiteboards, documents, drawings, blueprints—often neglect to include support for workspace awareness. This has contributed to their notable lack of success. Unlike the widespread use of communications systems such as e-mail and instant messaging, systems supporting a shared visual surface have not gained a broad following. This is surprising given that teams regularly work over shared workspaces in face-to-face settings.

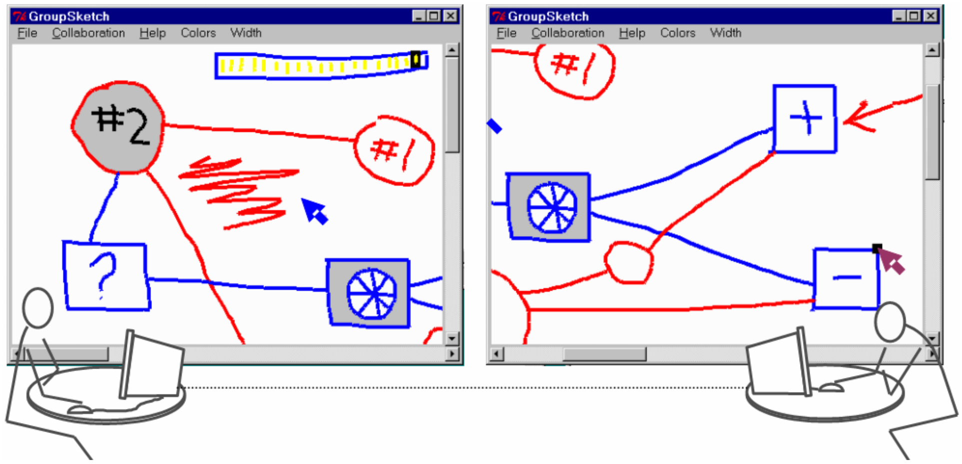

Figure 1. Sketch, a shared whiteboard

The problem is that maintaining this awareness has proved difficult in current real-time distributed systems in which information resources are poor and interaction mechanisms are foreign. Without good awareness, the ease and naturalness of collaboration is lost, making remote collaboration awkward, inefficient, and clumsy compared with face-to-face work. Thus, effective team cognition is compromised by the technology.

There are three main reasons why most groupware does not support workspace awareness. First, the input and output devices used in groupware systems generate only a fraction of the perceptual information that is available in a face-to-face workspace. Second, a user's interaction with a computational workspace generates much less information than actions in a physical workspace. Third, groupware systems often do not present even the limited awareness information that is available to the system.

As an example, consider the basic shared whiteboard in Figure 1 (Roseman & Greenberg, 1996). As each person draws, their actions are communicated to the other machine, so both participants' workspaces contain the same objects. At this moment in their task, the participants have scrolled their viewports to different parts of the workspace, and only a portion of their views overlap.

Systems such as this one show almost none of the awareness information that would be available to a co-located group working with a physical whiteboard.

People's hands and bodies are reduced to simple telepointers, there is no sound, and only a small piece of the entire drawing can be seen by a single person at one time. When different people scroll to different parts of the workspace (e.g., for pursuing individual activity), any information about where the other person is working or what they are doing is lost and can only be gathered through verbal communication. This system-imposed tunnel vision is equivalent to wearing blinders while working together.

Without this awareness, collaboration between team members in real time becomes awkward. It is difficult or impossible for the two participants to discuss particular objects, provide timely assistance, monitor the other person's activities, or anticipate that person's actions. Lack of information about others means that many of the little things that contribute to smooth and natural collaboration are missing from the interaction.

Chapter overview

In the remainder of this chapter, we argue that groupware designs and groupware systems must support workspace awareness. To do this, we first articulate the characteristics of workspace awareness typical in the everyday world, including what information people require, and the mechanisms they typically use to get it. This will help designers know what information must be captured, transmitted, and presented to all team members. Next, we introduce several interface techniques for actually capturing and presenting awareness information in our electronic workspaces. Finally, we validate the effectiveness of one interface technique by summarizing our experimental evaluations of it.

Awareness and workspace awareness

Awareness is knowledge created through interaction between an agent and its environment—in simple terms, "knowing what is going on" (Endsley 1995). Awareness has four basic characteristics (Adams, Tenney, & Pew, 1995; Endsley 1995; Norman 1993):

1. Awareness is knowledge about the state of a particular environment.
2. Environments change over time, so awareness must be kept up to date.
3. People maintain their awareness by interacting with the environment.
4. Awareness is usually a secondary goal—that is, the overall goal is not simply to maintain awareness but to complete some task in the environment.

Several types of awareness have been investigated in previous research. These include conversational awareness (e.g. Clark 1996), casual awareness of others in work groups (e.g. Borning & Travers 1991), and situation awareness (e.g. Gilson 1995, McNeese, Salas, & Endsley 2001). In particular, past work on *situation assessment* and *situation awareness* provides a wealth of theory and practice for our concept of workspace awareness. Specifically, situation assessment describes the human processes of gathering information (e.g., attention, pattern recognition, communication), wheras situation awareness is the end product resulting from effective situation assessment. Workspace awareness transforms and extends these ideas to the setting of a distributed visual workspace.

As mentioned in the introduction of this chapter, we define workspace awareness as *the up-to-the-moment understanding of another person's interaction with the shared workspace*. This definition bounds the concept in three ways. First, workspace awareness is an understanding of people in the workspace rather than just of the workspace itself. Second, workspace awareness is limited to events happening inside the workspace. Third, the physical nature of the workspace itself influences team cognition (which includes how people communicate and why they maintain workspace awareness): The combination of a working surface and the artifacts within it make the shared workspace both an external representation of the team's joint activity and its external memory (Clark 1996; Hutchins 1990; Norman, 1993). These constraints make workspace awareness a specialized kind of situation awareness, in which the situation comprises the other team members interacting with the workspace. The next sections describe in more detail a *framework for workspace awareness*. This

Category	Element	Specific questions
Who	Presence	Is anyone in the workspace?
	Identity	Who is participating? Who is that?
	Authorship	Who is doing that?
What	Action	What are they doing?
	Intention	What goal is that action part of?
	Artifact	What object are they working on?
Where	Location	Where are they working?
	Gaze	Where are they looking?
	View	How much can they see?
	Reach	How far can they reach?

Table 1. Elements of workspace awareness

framework articulates the elements of workspace awareness, the mechanisms by which it is maintained, and its uses in team cognition.

Workspace Awareness Framework

Part 1: What information makes up workspace awareness?

The first part of the framework divides the concept of workspace awareness into several elements of knowledge that answer basic "who, what, and where" questions about other team members and their activities. The elements reflect the fact that when we work with others in a physical shared space, we know who we are working with, what they are doing, and where they are working. Table 1 shows these elements and lists the questions that each element can answer. Note that the elements relate to awareness of present activities; Gutwin and Greenberg (2002) discuss additional elements that relate to the past.

"Who" awareness includes presence, identity and authorship (Table 1 top). Awareness of presence and identity is simply the knowledge that there are others in the workspace and who they are, and authorship involves the mapping between an action and the person carrying it out. "What" awareness covers actions, intentions and artifacts (Table 1 middle). Awareness of actions and intentions is the understanding of what another person is doing, either in detail or at a general level. Awareness of artifact means knowledge about what object a person is working on. “Where” awareness covers location, gaze, view and reach (Table 1 bottom). Location, gaze, and view relate to where the person is working, where they are looking, and what they can see. Awareness of reach involves understanding the area of the workspace where a person can change artifacts, because sometimes a person’s reach can exceed their view.

Although these elements appear to be merely common sense factors associated with one’s interaction with the environment, we suggest that they are the building blocks of team cognition and we further suggest that they form the foundation for robust, coordinated, and efficient team work in computer-supported cooperative work (CSCW). Before discussing the uses of workspace awareness, however, we first turn to the ways in which it is gathered in real-world settings.

Part 2: How is workspace awareness information gathered?

There are three main sources of workspace awareness information in face-to-face collaboration and three corresponding mechanisms that people use to gather it. People obtain information that is produced by people’s bodies in the workspace, from workspace artifacts, and from conversations and gestures. The mechanisms that they use to gather it are called consequential communication, feedthrough, and intentional communication.

Conversation, gesture, and intentional communication

A primary source of information that is ubiquitous in collaboration is conversation and gesture, and their mechanism is intentional communication (e.g. Birdwhistell, 1952; Clark 1996; Heath & Luff 1992). Verbal conversations are the prevalent form of communication in most groups, and there are three ways in which awareness information can be picked up from verbal exchanges. First, people may explicitly talk about awareness elements with their partners, and simply state where they are working and what they are doing. Explicit communication may also involve gestures and other visual actions (e.g. Short, Williams & Christie 1976).

Second, people can gather awareness information by overhearing others' conversations. Although a conversation between two people may not explicitly include a third person, it is understood that the exchange is public information that others can pick up. For example, Hutchins (1990) described how navigation teams on Navy ships talk on an open circuit, allowing everyone to hear each other's conversations, greatly adding to the team's resiliency in changing environments.

Third, people can pick up others' verbal shadowing, the running commentary that people commonly produce alongside their actions, spoken to no one in particular. Heath, Jirotka, Luff & Hindmarsh (1995) called this behaviour *outlouds* and suggested that they play a strong role in informing others about one's activities.

Bodies and consequential communication

Other important sources of awareness information in real-world collaboration are the other team members' bodies in the workspace. Because most activities that people do in a workspace are done through some visible bodily action, the position, posture, and movement of heads, arms, eyes, and hands "becomes an essential part of the flow of information fundamental for creating and sustaining teamwork" (Segal 1994 p. 24). Watching other people work is therefore a principal mechanism for gathering awareness information.

The mechanism of seeing and hearing other people active in the workspace is called *consequential communication*: information transfer that emerges as a consequence of a person's activity within an environment (Segal 1994). This kind of bodily communication is not intentional in the way that explicit gestures are: The producer of the information does not intentionally undertake actions to inform the other person, and the perceiver merely "picks up" what is available. Nevertheless, consequential communication provides a great deal of information.

Artifacts and feedthrough

The artifacts in the workspace are a third source of awareness information (e.g. Dix, Finlay, Abowd & Bale, 1993; Gaver 1991). By their positions, orientations, and movement, artifacts can show the state of people's interaction with them. Artifacts also contribute to the acoustic environment, making characteristic sounds when they are created, destroyed, or manipulated. Tools in particular have signature sounds, such as the snip of scissors or the scratch of a pencil.

The mechanism of determining a person's interactions through the sights and sounds of artifacts is called *feedthrough* (Dix et al., 1993). When artifacts are manipulated, they give off information, and what would normally be feedback to the person performing the action can also inform others who are watching or listening. When both the artifact and the actor can be seen, feedthrough is strongly coupled with consequential communication; at other times (such as in a groupware system) there may be a spatial or temporal separation between the artifact and the actor, leaving feedthrough as the only vehicle for information.

Part 3 – How do teams use workspace awareness?

Workspace awareness is used for many things in collaboration. Awareness can reduce effort, increase efficiency, and reduce errors for the activities of collaboration. This section describes three representative examples of activities that are aided by workspace awareness: management of coupling, simplification of verbal communication, and coordination of actions in the shared workspace.

Management of coupling

When people collaborate in a physical space, they shift seamlessly and effortlessly back and forth between individual and shared work (e.g. Dourish & Bellotti 1992; Gaver 1991). Salvador, Scholtz, and Larson (1996) called the degree to which people are working together *coupling*. Some of the reasons that people move from loose to tight coupling are that they see an opportunity to collaborate, that they need to discuss or decide something, that they need to plan the next activity, or that their current task requires another person's involvement.

Awareness of others' activities is crucial for smooth changes in coupling, both by helping people decide who they need to work wit, and by helping people decide when to make the transitions. This is verified in actual observations: For example Heath et al (1995) details how dealers in a financial office manage coupling by carefully monitoring their colleagues' activities.

Whether in an office or in a two-dimensional workspace, people try to keep track of others' activities when they are working in a loosely coupled manner, for the express purpose of determining appropriate times to initiate closer coupling. Without workspace awareness information, people will miss opportunities to collaborate or may interrupt the other person inappropriately.

Simplification of communication

Workspace awareness allows people to use the workspace and the artifacts in it to simplify their verbal communication, making team interaction more efficient. The type of communication we are interested in here is discussion involving task artifacts, which is a major part of the verbal activity in a shared workspace. In these conversations, the workspace can be used as a "conversational prop" (Brinck & Gomez 1992) – an external representation of the task that allows efficient nonverbal communication (Clark 1996, Hutchins 1990). Workspace awareness is important here because interpreting the visual signals depends on knowledge of where in the workspace they occur, what objects they relate to, and what the sender is doing. We illustrate the principle through three examples: deictic reference, visual evidence, and gaze awareness.

Deictic references. The practice of pointing or gesturing to indicate a noun used in conversation is called deictic reference, and is ubiquitous in shared workspaces (e.g. Seely Brown, Collins, & Duguid, 1989; Segal 1995; Tang 1991; Tatar, Foster, & Bobrow, 1991). Often, transcripts of verbal activity in a shared-workspace task cannot be correctly interpreted without a videotape of the workspace itself, because so many of the utterances contain words like "this one," "that one," "here," and "there" (e.g. Segal 1994). Deictic references allow communication to be much more efficient, primarily because constructing these 'indexical terms' without being able to point and gesture is very difficult.

Visual evidence. When people converse, they require evidence that their utterances have been understood. In verbal communication, a common form of this evidence is back-channel feedback. In shared workspaces, visual actions can also provide evidence of understanding or misunderstanding. Clark (1996, p. 326) provided an example from an everyday setting, in which Ben is getting Charlotte to center a candlestick in a display: "Okay, now, push it farther—farther—a little more—right there. Good.". Charlotte moves the candlestick after each of Ben's utterances, providing visual evidence that she has understood his instructions and has carried them out to the best of her interpretation.

Gaze awareness is knowing where another is looking and directing their attention (Ishii & Kobayashi 1992). It serves as visual evidence (to confirm that one is looking at the right place) and even as a deictic reference (as eye gaze can function as an implicit pointing act). It helps people monitor what others are doing. For example, if several people's gaze are directed at the same place, one can assume either that they are working together, or that one person is monitoring another person's actions.

The role of workspace awareness in deixis (i.e., where one's pointing or gesturing action disambiguates conversational references, such as when one says "this one" while pointing to an object), visual evidence and gaze awareness means that the elements of awareness are part of conversational common ground in shared spaces (Clark 1996). This implies that not only do you have to be aware of me to interpret my visual communication, but that I have know what you are aware of as well, so that I can safely make use of the workspace in my communication.

Coordination of actions

Coordinating actions in a collaborative activity means making them happen in the right order and at the right time to complete the task without conflicting with others in the group. Coordination can be accomplished in two ways in a shared workspace: “one is by explicit communication about how the work is to be performed…another is less explicit, mediated by the shared material used in the work process” (Robinson 1991). This second way is more efficient and much smoother, but requires that people maintain workspace awareness.

Awareness aids coordination, ranging from fine-grained continuous coordination on a tightly-coupled task, to coarse-grained occasional coordination on a loosely coupled task. Awareness eases coordination because it informs participants about the temporal and spatial boundaries of others’ actions, and because it helps them fit the next action into the stream. Workspace awareness is particularly evident in continuous action wherein people are working with the same objects. For example, CSCW researchers have noted that concurrency locks—in which the system monitors when one person selects an object and blocks others from using it—are less important or even unnecessary when participants have adequate information about what objects others are currently using (Greenberg & Marwood 1994). Another example is the way that people manage to avoid bumping into each others’ hands in a confined space. Workspace awareness allows people to track and predict others’ movements so as to coordinate access to the physical space or objects within it. Tang (1989) saw this kind of coordination when observing how small design groups managed their interaction over a tabletop: “the many ‘coordinated dances’ observed among the hands of the collaborators in the workspace … indicate a keen peripheral awareness of the other participants” (p. 95).

Many of the coordination characteristics that we think of in successful teams (“working like a well-oiled machine,” “singing off the same page”) mean, at least in artifact-based shared workspace, that the team is maintaining and using workspace awareness knowledge to track, predict, and mesh with the other members.

Summary of the framework

We defined workspace awareness as the up-to-the-moment understanding of another person's interaction with the shared workspace, and we introduced a three part workspace awareness framework. In the first part of our framework, the who, what, and where information elements are a starting point for thinking about the awareness requirements of particular task situations and provide a vocabulary for describing and comparing awareness support in groupware applications. The second part of the framework indicates how workspace awareness information is given off and gathered via intentional and consequential communication, as well as artifact state via feedthrough. The third part suggests that people actually use workspace awareness to manage their coupling and action coordination as they shift back and forth between individual and shared work. Now that we have discussed what workspace awareness is and how it works in collaboration, in the next section we turn to the issue of how it can be implemented in a groupware system to support distributed teams.

Supporting awareness in distributed groupware

The framework describes workspace awareness as it happens in face-to-face environments. When teams are distributed, however, it becomes much more difficult to maintain awareness of others, because groupware systems provide only a fraction of the information available in a physical workspace. In particular, two of the main awareness-gathering mechanisms – consequential communication and feedthrough – are greatly compromised in most systems.

In this section, we outline computational techniques that can be used to support workspace awareness in distributed groupware. We cover three main topics. First, we describe how *embodiments* can provide people with a representation in the workspace and provide a means for consequential communication. Second, we discuss the idea of *expressive artifacts* – workspace objects that maximize the amount of feedthrough information that is provided for the group's benefit. Third, we present *visibility techniques* that address the visibility problem in groupware, in which the narrow field of view prevents people from seeing others' awareness information that is situated in the workspace. We also give examples of these concepts to show how we can operationalize awareness support in our technology, which in turn supports effective distributed team cognition.

Embodiments

An embodiment is a visible representation that stands in for a person's body in a computational workspace. Embodiments are generally thought of as a way to provide a basic sense of presence in a virtual world, but they can also be a vehicle for both consequential and gestural communication. Although the limits of conventional input devices constrain an embodiment's expressiveness, they can still convey a great deal of awareness information.

There are three main types of embodiments used in distributed groupware: telepointers, avatars, and video images.

Telepointers

Telepointers are the simplest form of embodiment, and show the location of each team member's mouse cursor. Telepointers are effective at conveying awareness information, because the mouse cursor is the primary means by which people carry out actions in computational workspaces. In addition to simple cursor location, telepointers provide implicit information about presence, identity, activity, and even the specifics of an action.

In addition to the basic representation, telepointers can also be augmented to provide other awareness information (Greenberg, Gutwin, & Roseman, 1996). In the GroupSketch multi-user sketchpad, for example, telepointers are labeled with the name of their owner to show identity, change their shape to indicate what tool each person is using, and are oriented to different angles to distinguish one from the next (Greenberg & Bohnet 1991).

Avatars

Avatars are embodiments that represent people with stylized pictorial representations of actual bodies. They are primarily used in collaborative virtual environments where the world is shown in three dimensions (Benford, Fahlen, Greenhalgh, & Snowdon, 1995). Avatars provide a more humanlike body on which identity information and some kinds of gesture are more easily interpreted. The better avatars provide an embodiment that looks more or less like a person, have a recognizable face, and include an arm and hand to carry out any actions. Seeing the body in the workspace tells one where the distant avatar is located. Seeing the position of the face shows where it is looking. Seeing the hand shows where it is pointing (Fraser, Benford, Hindmarsh, & Heath, 1999).

Although avatars are an obvious choice in certain environments, the richer sense of presence that they provide does come at a cost: that the whole workspace be presented in a 3D or perspective view. Therefore, this technique must be weighed against the requirements for individual workspace interaction.

Video embodiment

Although video techniques go beyond the standard technical setup of most groupware systems, it is worth noting that several research systems have provided particularly effective video embodiments. These systems combine video images of team members with the representation of the computational workspace. What is particularly relevant is that the images are usually captured directly, and ideally one's body is in the "correct"' place relative to the workspace. This is important for correctly interpreting deictic references and gaze awareness information.

Video techniques provide a far more realistic and expressive embodiment than anything described above. There are several different ways that video can be used. First, with large display devices, silhouettes or shadows of people's bodies can be represented on the workspace (Tang & Minneman 1991). Second, full-fidelity video of arms and hands can provide detailed information about actions and movements (Tang & Minneman 1990). This allows a full range of motion (and two hands if needed) for gesturing over the artifacts in the workspace. Third, full-fidelity video of the entire upper body can show arms, hands, and faces (Ishii & Kobayashi, 1992), providing gaze awareness information and allowing eye contact. For example, Ishii's ClearBoard System (Ishii & Kobayashi, 1992) gives the impression of working with a remote collaborator through a pane of glass.

Expressive artifacts

Information produced by workspace artifacts – feedthrough – is one of the primary ways that people maintain workspace awareness. However, in computational workspaces, the interaction idioms and techniques used for manipulating artifacts often obscure people's actions, reducing feedthrough and compromising awareness. Unlike the physical world, interaction with computational environments is not limited to direct manipulation. *Symbolic manipulation* techniques are commands that let users specify actions in powerful and flexible ways. They are shortcuts using buttons, toolbars, and key commands that emphasize rapid invocation and execution. While often a good idea in single-user systems, symbolic manipulation produces minimal feedback (and thus minimal feedthrough), reducing people's ability to maintain awareness. This leads to three drawbacks for team members trying to stay aware of one another. First, symbolic actions have little or no visible representation in the workspace; actions are therefore harder to see in the workspace and are more likely to go unnoticed. Second, many symbolic actions are performed in similar ways so they are difficult to distinguish from one another. Third, symbolic actions can happen almost instantaneously, allowing little time for others to see and interpret them.

These problems can be addressed by transforming the minimal information provided by these actions to a more visible form as *feedthrough*. This approach makes artifacts more expressive: As Segal (1995) suggested, "compensate for consequential information that is lost…by providing enhanced feedback from the system indicating what specific actions each operator is performing" (p. 411). Below, we discuss two approaches—process feedthrough and action indicators—that make actions more obvious, distinguishable, and interpretable to others.

Process feedthrough

Some symbolic commands are invoked through interface widgets such as buttons, menus, or dialog boxes. The feedback provided from these command objects is never seen by other members of a distributed team: First, it is considered to be part of the application rather than part of the workspace, and second, it is considered to be distracting to other users. Feedback from these interfaces, however, can help the group to determine what actions people are composing. When other people receive this information, it becomes *process feedthrough*.

As a simple example of process feedthrough, consider a button in the interface of a groupware application. When a person's cursor moves over the button, it becomes highlighted on all users' screens; when a person presses the button, it is shown being pressed on all screens. While the highlight and the press give people a chance to interpret the action and determine what the other person is doing, it is very brief and easily missed. An alternative is to augment this natural feedthrough to make the action more visible. In Figure 2a, for example, the feedthrough of Carl's button press (left) is emphasized for Saul (right) by the remote button making a clicking sound and by adding a graphic that lingers longer than the actual press (Greenberg, Gutwin, & Roseman, 1996).

While buttons normally show too little feedthrough, menus may show too much –they carry a greater risk of distracting others or even obscuring their work. We can achieve better balance by muting the feedthrough. For example, Saul's menu in Figure 2b (right) displays only a portion of the feedback information that is visible to Carl, the local user (left).

Providing process feedthrough shows how actions are being composed and invoked but does not make the action itself more noticeable. When actions are hard to see, they can be augmented with artificial indicators, discussed next.

Action indicators and animations

Symbolic actions happen quickly and abruptly, making them hard to see and hard to interpret. For example, when someone presses the "delete" key to remove a selected object, the operation is nearly instantaneous. When actions are invisible, our approach is to create an artificial signal for them; these signals (called *action indicators*) can be given a more perceivable workspace representation.

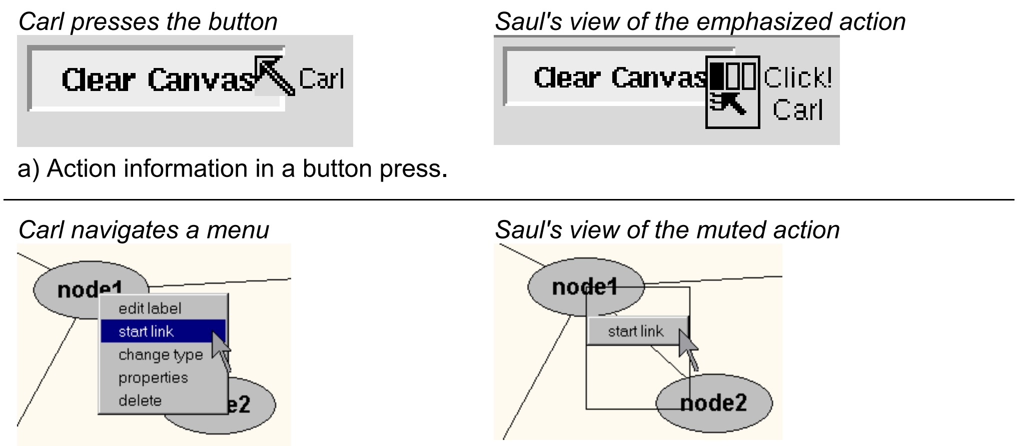

a) Action information in a button press.

b) a popup menu, as it is seen locally (left) and remotely (right).

Figure 2. Examples of Process Feedthrough

For example, an otherwise instantaneous delete operation can be made more obvious in several ways. One solution is to draw a text notification near the object on remote screens before removing the object, giving the rest of the group information and time to interpret its sudden disappearance. A more sophisticated solution, however, is to have the artifact itself visibly animate the otherwise invisible action, making it more perceptible. Figure 3 shows an example, in which the object labeled “node2” is deleted within a groupware concept map editor. The node does not simply disappear, but swells up for a moment (Frame 2) before gradually fading away (Frames 3 and 4). Through this “supernova”' effect, the effects of the delete action have been drawn out and made more noticeable.

A final indication technique uses sound cues. Sound has the advantage of being perceptible even when the object is off-screen, and can be combined with

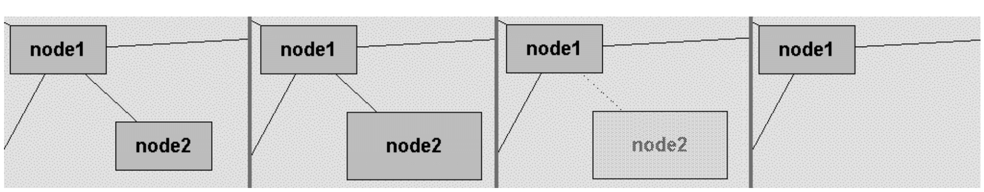

Figure 3. “Supernova” animation of a delete action.

the visual approaches described earlier. Different sounds can indicate different types of action, and can even convey characteristics and progress of the action (e.g., Gaver, 1991). For example, the system in Figure 3 plays a descending "whoosh" sound that fades away along with the visual representation of the deleted node.

Visibility techniques

Embodiments and expressive artifacts go a long way to restoring some of the workspace awareness information that is missing in a computational shared workspace. However, they are by nature situated in the workspace, that is, the information is produced at the workspace location where the action is taking place. This provides a valuable context for interpreting the information, but it also means that if a person is viewing a different part of the workspace, they will miss the information entirely. This is the visibility problem, and it occurs in groupware when the workspace is larger than the screen and when people can move their views independently.

There are a number of possible solutions to the visibility problem. The one we concentrate on here is the idea of providing multiple views of the workspace to give people different perspectives and greater visibility. In the following sections, we discuss three visibility displays: the radar view, the over-the-shoulder view, and the cursor's-eye view. Other possibilities described elsewhere include a variety of techniques that distort the workspace so that areas where others are working are larger than the areas that have no activity (Greenberg, Gutwin, & Cockburn, 1996).

Figure 4 illustrates and contrasts the various techniques. Figure 4a gives a bird's eye overview of the entire workspace, where two people Carl and Saul are working within it. Each can only see a portion of the workspace, as indicated by the bounding boxes. Figure 4b gives an example of what Saul may see. Most of his window shows a 'detailed view' where he sees a portion of the workspace at full size. The smaller add-on window at the upper left serves as a placeholder for

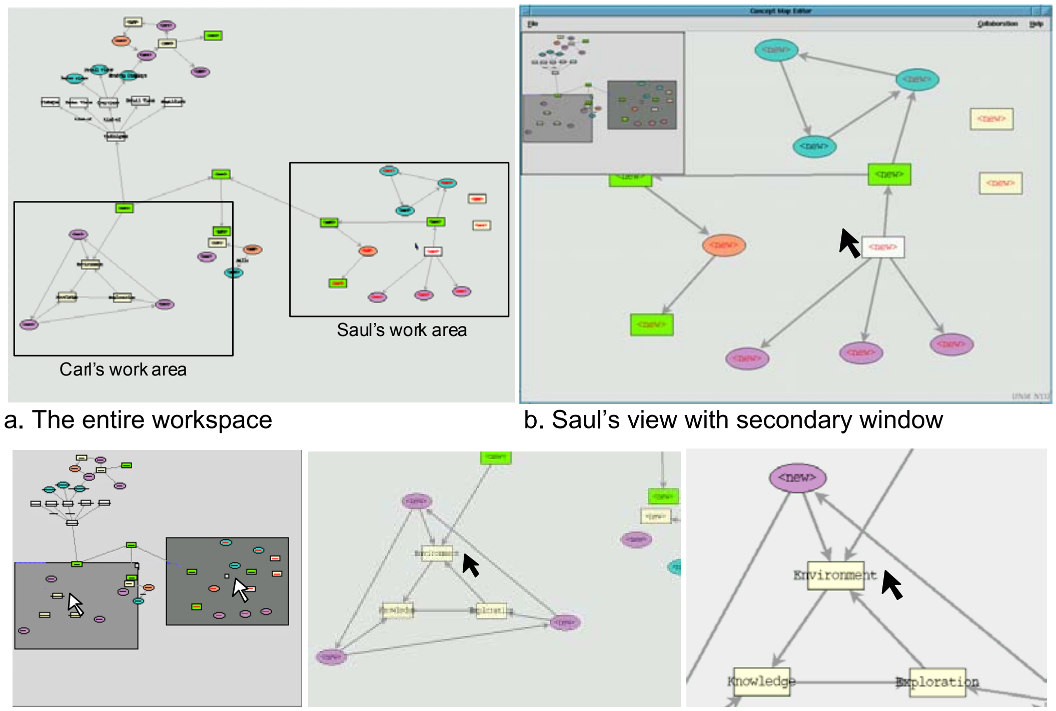

a. The entire workspace
b. Saul's view with secondary window
c. the radar view
d. Saul's over-the-shoulder view of Carl's area
e. cursor's-eye view of the area around Carl's cursor

Figure 4. Secondary views of the workspace for increased visibility.

awareness information, where it could take on one of the three forms shown in Figures 4c-e.

Radar views

Radar views are overview representations that show the entire workspace in miniature. They are usually presented as small windows inset into the main view. Although they do not take up much room, they provide a high-level perspective on artifacts and events in unseen areas of the workspace. For example, Figure 4b shows a radar view of a concept map workspace embedded as an inset window atop the detailed view; it is shown larger in Figure 5c. In the radar view, we see that two people's telepointers and main-view extents have been added to the basic overview. Augmentation such as this adds secondary embodiments to the display, and can therefore show who is in the space, where they are working (at two levels of detail), and what they are doing. However, because objects are shown at much lower resolution than in a normal view, radar views are best at helping people maintain high-level awareness of presence, locations, and general activities.

Over-the-shoulder views

The over-the-shoulder view shows a reduced version of another person's main view (Figure 4d), which is still larger than what would be seen in the radar view. The inspiration for this view is the idea of looking over at another person's work area in a face-to-face setting. This typically allows one to see what objects a coworker has in front of them, to see what they can see, and to look more closely at something that may have been noticed in peripheral vision. However, unlike the single radar view, separate over-the-shoulder views are needed to represent every participant.

Cursor's-eye views

A "cursor's-eye" view shows a small area directly around another person's mouse cursor (Figure 4e). Although its extents are limited, the cursor's-eye view shows objects and actions in full size and full detail. This view is useful when the precise details of another person's work are required. However, this view does not show the entire scene and, as with over-the-shoulder views, a separate view is required for each person.

Effects of Awareness Support on Groupware Usability

Of all of the awareness techniques and displays, the radar view is the one that we have found in practice to be the most useful in groupware applications (Gutwin, Roseman, & Greenberg 1996). To understand how the radar view affects group work in a measurable way, we conducted an experiment to test the effects of awareness support on groupware usability (Gutwin 1997; Gutwin & Greenberg 1999). We hypothesized that increased support for workspace awareness would improve the usability of groupware. In our study, we compared how distance-separated people performed various construction tasks using two awareness interfaces to a groupware pipeline construction system. The tasks embodied several essential activities that we believe are common to many workspace tasks. In particular,

- In the **Follow** task, one person goes to specific pre-defined locations to do some work on the pipeline. The other person, who does not know about these locations ahead of time, has to follow that person to those locations and do additional work to support the activities of the first person.
- In the **Copy** task, participants were asked to construct two identical structures. One person, the leader, had a picture of what was to be built. The other person, the copier, did not have the picture and so had to copy the leader's actions.
- In the **Direct** task, one participant was asked to verbally guide the other through six detailed construction tasks. The director had a map showing what was to be added and where but was not allowed to move around in the workspace.

The two groupware conditions differed only in the awareness information presented in the secondary awareness interfaces: the basic overview versus a radar view (Figure 5). Both were visible as insets on the main application (not shown, but similar to Figure 4b). The radar view and the overview differed in three ways, as compared in Figure 5.

1. *Update granularity*. The radar showed workspace objects as they moved, whereas the overview was only updated after the move was complete.
2. *Viewport visibility*. The radar showed both people's viewports (the area of the workspace visible in each person's main view), whereas the overview showed only the local user's viewport.
3. *Telepointer visibility*. The radar showed miniature telepointers for both users, whereas the overview did not show any telepointers.

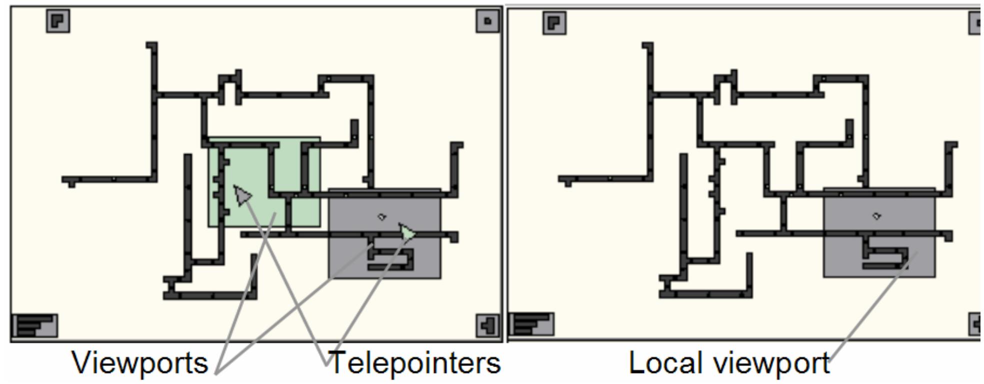

Figure 5. Radar view (left) and Overview (right).

A variety of results were obtained (see Gutwin & Greenberg 1999 for details). When using the radar view, groups finished the Follow and Direct tasks significantly faster (about 3 minutes with the radar and about 4.5 with the overview). Also, groups using the radar view spoke significantly fewer words in the Follow task (about 100 words with the radar view and about 225 with the overview). No differences were found in perceived effort for any of the tasks, and no differences were found on any measure for the Copy task. After all tasks were completed and pairs had used both interfaces, participants were asked which system the participant preferred overall. All of the 38 people who responded chose the radar view over the basic overview.

The primary reasons for the radar view's success is that visual awareness information makes it easier to communicate useful information without talking; and awareness information gives people confirmation about the other person's activities. Thus, our data support the aspect of our awareness framework highlighting the criticality of both visual and conversational common ground in facilitating coordination. For example, we saw that the radar condition provided visual indication of the other person's location and activity by showing view rectangles and telepointers. This information helped people complete the Follow and Direct tasks more quickly. In the Follow task, for example, followers could simply watch where the leader's view rectangle went on the screen and then go there themselves; in contrast, the overview condition forced people to construct complicated verbal directions to tell the other person where to go. The radar view transformed the task from a verbal one to a visual one, making it simpler and more efficient. This transformation also explains why groups used significantly fewer words in the Follow task when they used the radar view.

We also saw that the radar view provided continuous feedthrough about location and object position, which allowed groups to complete the Follow and Direct tasks more quickly. In particular, this feedback gave people visual evidence of understanding (Brennan, 1990), which was more effective and less error-prone than verbal evidence. This difference was particularly apparent in the Direct task, in which the director guides the actor's movement by giving her an instruction. With each instruction, the director requires evidence that he has succeeded in conveying the correct meaning to the actor, and that the actor has successfully moved where she is supposed to go. In addition, the director often cannot give the next instruction until he knows that the actor has successfully completed the current one. The information differences between the radar view and the overview provided directors with different kinds of evidence, and afforded different means – verbal versus visual – for establishing that instructions had been understood and carried out. In the overview, actors had to verbally acknowledge that they had completed the direction (e.g. "OK, I'm there"); this confirmation, however, is given at the end of the action, and if the action has been in error, considerable effort has been wasted while the actor went the wrong way. In contrast, the radar view showed up-to-the-moment object movement and viewport location. In the Direct task, these representations could be used as immediate visual evidence of the actor's understanding and intentions. If the actor started moving the wrong way, the director would see the misunderstanding immediately and could interrupt the actor to correct the action. In addition, the availability of continuous evidence made it possible for people to give continuous instructions. This is a strategy with far fewer verbal turns, and where the actor acknowledges implicitly through his or her actions. Clark (1996) summarized the difference between verbal and visual acknowledgment for on-going "installment" utterances like instructions: "in installment utterances, speakers seek acknowledgments of understanding (e.g. 'yeah') after each installment and formulate the next installment contingent on that acknowledgment. With visual

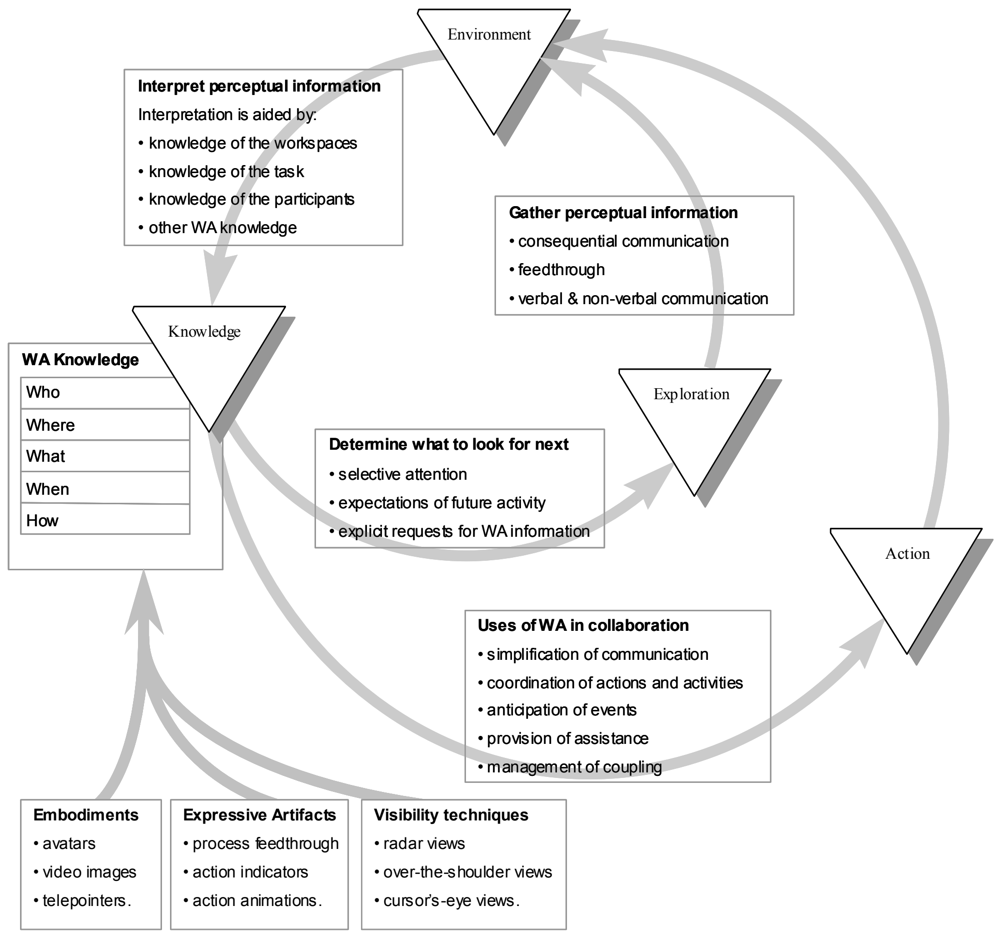

Figure 6. Workspace awareness framework and techniques

evidence, [the speaker] gets confirmation or disconfirmation while he is producing the current installment" (p. 326).

In summary, evidence of understanding and action in the radar was accurate, easy to get, and timely. The director was able to determine more quickly whether the instruction was going to succeed and could reduce the cost of errors. These results add weight to the overall hypothesis that awareness support improves groupware usability.

Summary of awareness support

We believe these various computational techniques—embodiments, expressive artifacts, and visibility techniques—have promise for distributed team cognition because they operationalize workspaces awareness support in technology.

Of all of the awareness techniques and displays, the radar view is the one that we have found to be the most useful in groupware applications. Our investigations comparing awareness techniques show that the basic overview itself is valuable when the workspace is larger than the screen, and that the feedthrough and consequential communication provided in the radar view allow people to maintain workspace awareness even when their collaborators are out of view (Gutwin, Roseman, & Greenberg 1996).

Conclusions

Our main message is: that for people to sustain effective team cognition when working over a shared visual workspace, groupware systems must give team members a sense of workspace awareness.

In this chapter, we have explored several issues that must be considered before this message can be implemented effectively. These issues are illustrated in the diagram of Figure 6. First, designers need a better understanding of what exactly is meant by workspace awareness. This is the role of our workspace awareness framework, in which we described what information makes up workspace awareness, how workspace awareness information is gathered, and how teams use it. Second, developers need a repository of computational interaction techniques that support workspace awareness if they are to codify it within actual systems. We described several such techniques, including various forms of embodiments that give off bodily expressions, expressive artifacts for showing feedthrough, and three visibility techniques for displaying awareness information when people are looking at different parts of a workspace. Third, we need to show that these techniques are effective. As an example, we summarized a study that we have done that looks at the fine-grained effects of several awareness techniques and that validates where they are useful.

Unlike the everyday world in which awareness just "happens" as teams work within it, designers of distributed shared workspace groupware must explicitly program in features to gather awareness information, to transmit that information down the communication channel, and to display it effectively on the screen. This will only happen if we give designers a good understanding of workspace awareness and a proven repertory of interaction techniques that support it.

Acknowledgments

We are grateful to Microsoft Research, Intel Corporation, the Alberta Software and Engineering Research Consortium (ASERC), the National Sciences and Engineering Research Council of Canada (NSERC) and Smart Technologies who provided funding, some equipment, and encouragement at various phases of this work.

References

Adams, M., Tenney, Y., & Pew, R. (1995) Situation Awareness and the Cognitive Management of Complex Systems, *Human Factors*, 37(1), 85-104.

Benford, S., Bowers, J., Fahlen, L., Greenhalgh, C., & Snowdon, D. (1995) User Embodiment in Collaborative Virtual Environments, *Proceedings of the Conference on Human Factors in Computing Systems* (CHI'95), 242-249, New York: ACM.

Birdwhistell, Ray L. (1952*) Introduction to kinesics : an annotation system for analysis of body motion and gesture*. Louisville, KY: University of Kentucky Press.

Borning, A., & Travers, M. (1991) Two Approaches to Casual Interaction over Computer and Video Networks, *Proceedings of the Conference on Human Factors in Computing Systems*, 13-19, New York: ACM.

Brennan, S. (1990) *Seeking and Providing Evidence for Mutual Understanding*, Unpublished doctoral dissertation, Stanford University, Stanford, CA.

Brinck, T., & Gomez, L. M. (1992) A collaborative medium for the support of conversational props, Proceedings of *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (CSCW'92), 171-178, New York: ACM.

Cannon-Bowers, J, Salas, E., & Converse, S. (1993). Shared mental models in expert decisionmaking teams. In N. J. Castellan, Jr. (Ed.), *Current issues in individual and group decision making* (pp. 221-246). Hillsdale, NJ: Erlbaum.

Clark, H. (1996) *Using Language*, Cambridge, England: Cambridge University Press.

Dix, A., Finlay, J., Abowd, G., & Beale, R. (1993) *Human-Computer Interaction*, New York: Prentice Hall.

Dourish, P., & Bellotti, V. (1992) Awareness and Coordination in Shared Workspaces, *Proceedings of the Conference on Computer-Supported Cooperative Work*, 1992, 107-114, New York: ACM.

Ellis, C., Gibbs, S., & Rein, G. (1991) Groupware: Some Issues and Experiences, *Communications of the ACM*, 34(1), 38-58.

Endsley, M. (1995) Toward a Theory of Situation Awareness in Dynamic Systems, *Human Factors*, 37(1), 32-64.

Fraser, M., Benford, S., Hindmarsh, J., & Heath, C. (1999) Supporting awareness and interaction through collaborative virtual interfaces. *Proceeding of ACM UIST99*, 27-36, New York: ACM.

Gaver, W. (1991) Sound Support for Collaboration, *Proceedings of the Second European Conference on Computer Supported Cooperative Work*, 293-308, Amsterdam: Kluwer.

Gilson, R. D. (1995) Introduction to the Special Issue on Situation Awareness, *Human Factors*, 37(1), 3-4.

Greenberg, S., & Bohnet, R. (1991) GroupSketch: A multi-user sketchpad for geographically-distributed small groups, *Proceedings of Proceedings of Graphics Interface '91*, Calgary, Alberta, 207-215, San Francisco: Morgan-Kaufmann.

Greenberg, S., Gutwin, C. and Cockburn, A. (1996) Using Distortion-Oriented Displays to Support Workspace Awareness. in *People and Computers XI (Proceedings of the HCI'96),* A. Sasse, R.J. Cunningham, and R. Winder, Editors. 299-314, New York: Springer-Verlag.

Greenberg, S., Gutwin, C., & Roseman, M. (1996) Semantic Telepointers for Groupware. *Proceedings of the OzCHI '96 Sixth Australian Conference on Computer-Human Interaction*. 54-61, Los Alamitos, CA: IEEE Computer Society.

Greenberg, S., & Marwood, D. (1994) Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface, *Proceedings of the Conference on Computer-Supported Cooperative Work*, 207-217, New York: ACM.

Gutwin, C., & Greenberg, S. (1999) Effects of Awareness Support on Groupware Usability. *ACM Transactions on CHI*, vol. 6 no. 2, 243-281.

Gutwin, C. (1997) *Workspace Awareness in Real-Time Distributed Groupware*. Unpublished Doctoral Dissertation, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada.

Gutwin, C., & Greenberg, S. (2002) Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work*, 11(3-4), 2002, Kluwer Academic Press.

Gutwin, C., Roseman, M. & Greenberg, S. (1996) A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. *Proceedings of ACM CSCW'96 Conference on Supported Cooperative Work*, 1996, New York: ACM.

Heath, C., Jirotka, M., Luff, P., & Hindmarsh, J. (1995) Unpacking Collaboration: the Interactional Organisation of Trading in a City Dealing Room, *Computer Supported Cooperative Work*, 3(2), 147-165, 1995, Kluwer Academic Press.

Heath, C., & Luff, P. (1992) Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms., *Computer-Supported Cooperative Work*, 1(1-2), 69-94, 1992, Kluwer Academic Press.

Hutchins, E. (1990) The Technology of Team Navigation, in *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, J. Galegher, R. Kraut and C. Egido ed., 191-220, Hillsdale, NJ: Lawrence Erlbaum.

Hutchins, E. (1996) *Cognition in the wild*. Cambridge, MA: MIT Press.

Ishii, H., & Kobayashi, M. (1992) ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact, *Proceedings of the Conference on Human Factors in Computing Systems*, 525-532, New York: ACM.

McGrath, J. (1984) *Groups: Interaction and Performance*, Englewood Cliffs, NJ: Prentice-Hall.

McNeese, M., Salas, E., & Endsley, M. (2001) *New Trends in Cooperative Activities: Understanding System Dynamics in Complex Environments*. San Diego, CA: Human Factors and Ergonomics Society Press.

Norman, D. (1993) *Things That Make Us Smart*, Reading, Mass:Addison-Wesley.

Robinson, M. (1991) Computer-Supported Cooperative Work: Cases and Concepts, *Proceedings of Groupware '91*, 59-75, Utrecht, Netherlands: Software Engineering Research Centre.

Roseman, M., & Greenberg, S. (1996) Building Real-Time Groupware with GroupKit, a Groupware Toolkit, *ACM Transactions on Computer-Human Interaction*, 3(1), 66-106.

Salvador, T., Scholtz, J., & Larson, J. (1996) The Denver Model for Groupware Design, *SIGCHI Bulletin*, 28(1), 52-58.

Seely Brown, J., Collins, A., & Duguid, P. (1989) Situated Cognition and the Culture of Learning, *Educational Researcher* (January-February), 32-42.

Segal, L. (1994) *Effects of Checklist Interface on Non-Verbal Crew Communications*, NASA Ames Research Center, Contractor Report 177639, Moffett Field, CA: NASA.

Segal, L. (1995) Designing Team Workstations: The Choreography of Teamwork, in *Local Applications of the Ecological Approach to Human-Machine Systems*, P. Hancock, J. Flach, J. Caird and K. Vicente ed., 392-415, Hillsdale, NJ:Lawrence Erlbaum.

Short, J., Williams, E., & Christie, B. (1976) Communication Modes and Task Performance, in Readings in *Groupware and Computer Supported Cooperative Work: Assisting Human-Human Collaboration*, R. M. Baecker ed., 169-176, Mountain View, CA:Morgan-Kaufmann Publishers.

Tang, J. (1989) *Listing, Drawing, and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams*, Unpublished doctoral dissertation, Stanford University, Stanford, CA.

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

Tang, J. C., & Minneman, S. L. (1990) Videodraw: A video interface for collaborative drawing, *Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems*, 313-320, New York: ACM.

Tang, J. C., & Minneman, S. L. (1991) VideoWhiteboard: Video shadows to support remote collaboration, *Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems*, 315-322, New York: ACM.

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