The Importance of Awareness for Team Cognition in Distributed Collaboration

Carl Gutwin¹ and Saul Greenberg²

¹University of Saskatchewan, Saskatoon, Saskatchewan, S7N 5A9, Canada

²University 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 things 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 things like using environmental cues to establish a common ground of understanding, seeing who is around and what they are doing, monitoring the state of artefacts 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 will 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*.

Before getting into details, we will set the scene by first describing the collaborative situations we address in this chapter, and then by introducing workspace awareness and why it is a problem in conventional groupware systems.

Cite as:

Gutwin, C., and Greenberg, S. (In submission) **The importance of awareness for team cognition in distributed collaboration**. In E. Salas, S. M. Fiore and J. A. Cannon-Bowers (Editors) *Team Cognition: Process and Performance at the Inter- and Intra-individual Level*. APA Press.

The collaborative situations we address

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 et al. 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) where 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, where people shift frequently between individual and shared activities during a work session (e.g. Dourish and Bellotti, 1992; Salvador et al., 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 these shared workspaces. While we recognize that certain task domains may require people to follow an explicit process, peoples' actions as they perform rudimentary workspace operations are typically graceful and unconscious. Similarly, the team maintains a shared mental model as each member synergistically tracks the natural 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 and Greenberg, in press).

We will elaborate on this definition in later sections, but for now we will just say that workspace awareness is limited to those things 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 we 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 email 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.

The problem is that maintaining this awareness has proven difficult in current real-time distributed systems where 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. This system is included as demonstration software within the GroupKit groupware toolkit (Roseman and 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.

Figure 1. Sketchpad, a shared whiteboard

Systems like 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. In the situation pictured in Figure 1, it will be difficult or impossible for the two participants to discuss particular objects, provide timely assistance, monitor the other person's activities, or anticipate their actions. In short, lack of information about others means that many of the little things that contribute to smooth and natural collaboration will be missing from the interaction.

What is in this chapter

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: 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 and Pew 1995; Norman 1993; Endsley 1995):

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.

Adams et al (1995) suggest a cognitive model that shows how awareness is maintained in dynamic environments, a model that also draws together both the process and product aspects of awareness. The model is Neisser’s (1976) perception-action cycle. Neisser’s model captures the interaction between the agent and the environment, and incorporates relationships between a person’s knowledge and their information-gathering activity.

Several types of awareness that have been investigated in previous research, including conversational awareness (e.g. Clark 1996), casual awareness of others in work groups (e.g. Borning and Travers 1991), and situation awareness (e.g. Gilson 1995). In particular, past work on situation awareness provides us with a starting point for our conception of workspace awareness.

As mentioned in the introduction, 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 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, Norman 1993, and Hutchins

1990). These constraints make workspace awareness a specialized kind of situation awareness – where the situation comprises the other team members interacting with the workspace. The next sections describe in more detail a *framework for workspace awareness*. As we will see, this framework articulates the elements of workspace awareness, the mechanisms by which it is maintained, and its uses in team cognition.

The 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. Discussion of additional elements that relate to the past can be found in (Gutwin and Greenberg, in press).

‘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 things, since sometimes a person’s reach can exceed their view.

The elements of workspace awareness are all commonsense things that deal with interactions between a person and the environment, but as we will show, this information is extremely valuable in helping groups to be robust, well-coordinated, and efficient. Before discussing the uses of workspace awareness, however, we first turn to the ways in which it is gathered in real-world settings.

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

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. Clark 1996; Heath and Luff 1992; Birdwhistell, 1952). 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, and 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

others' 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, Jirotko, Luff and Hindmarsh (1995) have observed this behaviour, which they call "outlouds." They note that although these "outlouds...might be thought relatively incursive, potentially interrupting activities being undertaken by [others] in the room, [they are] perhaps less obtrusive than actually informing particular persons" (p. 157).

Bodies and consequential communication

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

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, however, 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. For example, in a study of piloting teams, Segal reports that:

[Pilots] spent most of their time—over 60%—looking across at their [partner's] display while it was being manipulated. This suggests that beyond the information provided by the display itself, these pilots were specifically looking for information provided by the dynamic interaction between their crewmembers and that display (p. 24).

Artifacts and feedthrough

The artifacts in the workspace are a third source of awareness information (e.g. Dix et al 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 and Bellotti 1992; Gaver 1991). Salvador et al (1996) call 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 with, and by helping people decide when to make the transitions. Heath et al (1995) give an example of the latter. In a financial office, dealers manage coupling by carefully monitoring their colleagues' activities:

...though dealers may be engaged in an individual task, they remain sensitive to the conduct of colleagues and the possibility of collaboration... 'Peripheral' monitoring or participation is an essential feature of both individual and collaborative work within these environments. ...So, for example, it is not unusual in the dealing room for individuals to time, with precision, an utterance which engenders collaboration, so that it coincides with a colleague finishing writing out a ticket or swallowing a mouthful of lunch. By monitoring the course of action in this way and by prospectively identifying its upcoming boundaries, individuals can successfully initiate collaboration so that it does not interrupt an activity in which a colleague is engaged. (p. 152)

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 and Gomez 1992) – an external representation of the task that allows efficient nonverbal communication (Hutchins 1990; Clark 1996). 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. Segal 1995; Tatar et al 1991; Tang 1991). Often, transcripts of verbal activity in a shared-workspace task cannot be correctly interpreted without a videotape of the workspace itself, since 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. As Seely Brown, Collins and Duguid (1989) state: “the best way to discover the importance and efficiency of indexical terms and their embedding context is to imagine discourse without them. Authors of a collaborative work will recognize the problem if they have ever discussed the paper over the phone. ‘What you say here’ is not a very useful remark. *Here* in this setting needs an elaborate description (such as ‘page 3, second full paragraph, fifth sentence, beginning...’)” (p. 36). Team members are freed from these complex utterances if they can maintain awareness of pointing fingers and the artifacts that they are pointing at in the workspace.

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) provides an example from an everyday setting, where Ben is getting Charlotte to center a candlestick in a display – “Okay, now, push it farther—farther—a little

more—right there. Good.” (p. 326). 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. This kind of evidence can be used whenever people carry out joint projects involving the artifacts in a shared workspace.

Gaze awareness is knowing where another is looking (Ishii and Kobayashi 1992). It helps one know where another is directing their attention. Thus it also 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 a person’s gaze is directed at a portion of the workspace where no-one is working, one can assume that they are pursuing individual work. If several people’s gaze are directed at the same place, one can assume that they are either working together, or that one person is monitoring another person’s actions.

The role of workspace awareness in deixis, 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, p. 42). This second way is more efficient and much smoother, but requires that people maintain workspace awareness.

Awareness aids both fine and coarse-grained coordination, since it informs participants about the temporal and spatial boundaries of others’ actions, and since it helps them fit the next action into the stream. Workspace awareness is particularly evident in continuous action where people are working with the same objects. For example, CSCW researchers have noted that concurrency locks are less important or even unnecessary when participants have adequate information about what objects others are currently using (Greenberg and 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 in design activity:

the physical closeness among the participants...allows a peripheral awareness of the other participants and their actions, as evidenced in the many 'coordinated dances' observed among the hands of the collaborators in the workspace. There were many episodes of intricate coordinated hand motions, such as getting out of the way of an approaching hand or avoiding collisions with other hands. These coordinated actions 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 is the up-to-the-moment understanding of another person's interaction with the shared workspace. We introduced a three-part workspace awareness framework that articulates: its component elements, the mechanisms used to maintain it, and its uses in collaboration.

In the first part, the component elements are the information that makes up workspace awareness, the 'who, what, and where' questions. They deal with issues like who is present and who is responsible for actions, where people are working and where they can see, and what actions they are performing and what their intentions are. The 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. Intentional communication includes verbal conversation, visible gestures, overhearing conversations, and verbal shadowing. Consequential communication is information that emerges as a consequence of a person's bodily activities within an environment. Artifacts and feedthrough show the state of artifacts and how people are interacting with them.

The third part suggests ways that people actually use workspace awareness information as part of their collaboration. People use it to manage their coupling, as they shift back and forth between individual and shared work. Good awareness simplifies communication as the information it provides means that many things do not have to be negotiated or stated explicitly. It helps people coordinate their actions by helping them happen in the right order and at the right time without conflict.

Now that we have discussed what workspace awareness is and how it works in collaboration, we turn in the next section 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 small fraction of the information that is 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, where the narrow field of view prevents people from seeing others’ awareness information that is situated in the workspace.

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, since 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, and Roseman, M. 1996). For example, telepointers can be colored or tagged to show identity, and can change shape to indicate what tool each person is using. For example, the telepointers in Groupsketch (Greenberg and Bohnet 1991) were annotated with the participant's name, were oriented to different angles, and changed to a different tool icon depending on whether the user was drawing, erasing, typing, or pointing (Figure 2).

Avatars

Avatars are embodiments that represent people with stylized pictorial representations of actual bodies. They are primarily used in collaborative virtual environments (CVEs) where the world is shown in three dimensions (Benford et al, 1995). Avatars provide a more humanlike body on which identity information and some kinds of gesture are more easily interpreted. So, instead of a telepointer with a name tag, avatars provide an embodiment that looks more or less like a person, can have a recognizable face, and whose actions in the workspace are carried out by a "hand" at the end of an "arm." Figure 3 illustrates one example, where two avatars are present in the space, and we can see where the distant avatar is located in the workspace, where it is looking, and where it is pointing.

Figure 2. Groupsketch, showing modal telepointers
(Greenberg et al 1991).

Figure 3. Avatars in the MASSIVE virtual environment (Fraser et al, 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 embodiments through video. 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, where 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 a 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 (e.g. Tang and Minneman 1991). Second,

full-fidelity video of arms and hands can provide detailed information about actions and movements (e.g. Tang and 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 (e.g. Ishii et al 1992), providing gaze awareness information and allowing eye contact. For example, Ishii's ClearBoard System (Ishii et al 1992), illustrated in Figure 4, gives the impression of working with a remote collaborator through a pane of glass.

Figure 4. ClearBoard, showing video image of remote user (Ishii et al 1992)

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 that are 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, however, by transforming the minimal information provided by these actions to a more visible form as feedthrough. This is the approach of making artifacts more expressive: as Segal (1995) suggests, "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 can make actions more obvious, more distinguishable, and more 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. The highlight and the press give people a chance to interpret the action and determine what the other person is doing. If the button represents a particularly important action, the natural feedthrough can even be augmented to make it more visible. For example, the button could make a clicking sound when it is pressed, or use a more obvious highlight colour, or can enhance the action itself (Figure 5).

a) Carl presses the button

b) Saul's view of it

Figure 5. Action information in a button press (Greenberg, Gutwin, and Roseman, 1996).

A second example involves process feedthrough for menus. Menus are a more complicated case than buttons, since they carry a greater risk of distracting others or even obscuring their work. To reduce this risk, we display only a portion of the feedback information that is visible to the local user. Figure 6 shows a group-visible popup menu as it appears on a local and a remote display. This technique is useful when a large menu would hide too much of the local user's view.

Figure 6. A popup menu, as it is seen locally (left) and remotely (right).

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, an approach we discuss 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.

For example, a delete operation can be made more obvious in several ways. One simple solution is to draw a text notification near the object on remote screens before removing the object. This technique gives the rest of the group information and time to interpret the sudden disappearance of the object. A more sophisticated solution, however, is to have the artifact itself animate the action. When actions

cause a visible change in the artifact, these changes can be made more perceptible even if the action is invisible. Figure 7 shows this approach in a concept-map system. When the object labeled ‘node2’ is deleted, it does not simply disappear, but swells up for a moment before gradually fading away (the supernova effect). Although the original delete action is still invisible, the effects of that action have been drawn out and made more noticeable.

Figure 7. “Supernova” animation of a delete action.

A final indication technique uses sound cues to indicate actions. Sound has the advantage of being perceptible even when the object is off-screen, and can be combined with the visual approaches described above. 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 shown in Figure 7 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 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 paragraphs, 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, and Cockburn 1996).

We will use Figure 8 as an example to illustrate and contrast various techniques. Figure 8a give 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 8b 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 awareness information, where it could take on one of the three forms shown in Figures 8c-e.

Figure 8. Secondary views of the workspace for increased visibility.

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. In addition, radar views can make workspace awareness information visible, regardless of where in the workspace it is situated.

For example, Figure 8b shows a radar view of a concept map workspace embedded as an inset window atop the detailed view; it is shown larger in Figure 8c. In the radar view, we see that two people's telepointers and main-view extents have been added to the basic overview. When the radar view is augmented like this, it essentially 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, since 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. The objects are shown much smaller than full size, but are still considerably larger than they would be in an overview. This display has enough resolution to represent actual object manipulations and therefore provides a more detailed sense of awareness about activity. The inspiration for this view is the idea of looking over at another person's work area in a face-to-face setting, to see what objects they have 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.

Figure 8d shows an over-the-shoulder view of Carl's work area. By adding this view to his interface (i.e., by replacing the inset in Figure 8b), Saul can keep track of exactly what Carl can see, and can tell for most purposes what Carl is doing in his part of the workspace. However, unlike radar views, multiple participants means that each person's screen needs to display multiple over-the-shoulder views, i.e., one extra inset for each additional person.

Cursor's-eye views

A "cursor's-eye" view shows a small area directly around another person's mouse cursor. 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: for example, when one person wishes to keep an eye on the way another aligns objects, types, or adjusts the fine details of a drawing. The cursor's-eye view does not show the entire scene, so its use is limited to situations where the general nature of a person's actions are already known. A cursor's-eye view of Carl's cursor is shown in Figure 8e. As with over-the-shoulder views, each extra participant requires an extra inset on the display.

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. An investigation we carried out to compare several awareness techniques showed 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 allow people to maintain workspace awareness even when their collaborators are out of view (Gutwin, Roseman, and Greenberg 1996). To determine if this display could affect group work in a measurable way, we carried out an experiment to test the effects of awareness support on groupware usability. This experiment is described in the next section.

The Effects of Awareness Support on Groupware Usability

Based on our experiences with building and using awareness displays in real-time distributed groupware systems, we hypothesized that increased support for workspace awareness would improve the usability of groupware. In the study, we compared two groupware interfaces that provide different amounts of awareness information through their overview displays. In particular, we compared a basic overview to a radar view that showed viewport locations, miniature telepointers, and object motion as the objects were manipulated.

We summarize the study highlights here. A detailed description of both the methodology, its results and our interpretation of them are given in (Gutwin 1997; Gutwin and Greenberg 1999).

System and experimental conditions

The experimental application was a pipeline construction kit that allows the assembly and manipulation of simple pipeline networks in a shared two-dimensional workspace (Figure 9). Users can create, move, and rotate sections of pipe, and can join or split sections using a welding tool. The workspace is rectangular, and four times larger than the computer screen in each direction. Users scroll around the workspace by dragging their cursor past the window border.

The pipeline system's interface consists of two windows. The main view allows users to manipulate objects and to scroll to other areas of the workspace. People create pipelines by dragging pipe sections from storehouses in the corners of the workspace (see Figure 9), aligning the sections, and then welding them together by dropping a diamond-shaped welding tool onto the joint. Welds are marked by a yellow square, and once pieces are welded, they move as a unit.

Figure 9. The pipeline application (radar view version)

The second window is one of two miniature views – either the *radar view* or the *overview*. This view is inset into the top left corner of the main view, and shows the entire workspace in miniature. The radar view and the overview differed in three ways, as compared in Figure 10.

1. *Update granularity*. The radar showed workspace objects as they moved; 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) and the overview showed only the local user's viewport.
3. *Telepointer visibility*. The radar showed miniature telepointers for both users, and the overview did not show any telepointers.

Figure 10. Radar view (left) and Overview (right).

In sum, the two conditions differed only in the awareness information presented in the miniature. The overview only showed information about the local user, while the radar showed where the other person was located, showed their pointer, and showed moves as they occurred.

Note that this study only looks at fine-grained aspects of how groupware can provide awareness information. In a previous usability study, it was abundantly clear that systems that provide some awareness information are far better than those that provide no awareness information (Gutwin, Roseman, and Greenberg, 1996).

Tasks

Participants completed three different tasks. The tasks were designed from episodes of joint actions that we had previously seen in face-to-face collaboration, episodes that required people to move independently around the workspace, and required that people maintain a sense of workspace awareness.

The **Follow** task involved meeting another person at a specified location in the workspace. Participants were asked to make ten specific welds on an existing pipe network. One person, the joiner, was given a map showing the locations to be welded, and had to prepare the pipe sections at each place. The other person was the welder, and would follow the joiner to each location and weld the pipe. Since the welder had no map, the joiner was also responsible for ensuring that the welder went to the correct location.

The **Copy** task involves indicating objects to another person. Participants were asked to construct two identical structures from two existing stockpiles of pipe sections. The stockpiles were located at opposite ends of the workspace. One

person, the leader, had a picture of what was to be built, and used this to find the next piece in their stockpile. The other person, the copier, did not have the picture, and so had to copy the leader's actions. The leader was responsible for making sure that the copier knew which piece to take next and where to place it.

The **Direct** task involves giving workspace directions. One participant was asked to verbally guide the other through adding six specific pipe sections to an existing network. The director had a map showing which pieces were to be added, and where they were to be added, but was not allowed to move around in the workspace. The actor did the work, following the director's instructions. The director used only the miniature view for this task.

Study design

The design combines two independent variables in a two-way mixed factorial design: View (overview or radar) is a between-participants factor, and Task (follow, copy, or direct) is a repeated-measures factor. The specific hypotheses were:

1. Groups in the radar condition will complete the first three tasks more quickly, with greater efficiency, and with greater satisfaction.
2. Groups who use the overview first and then the radar view will have a greater improvement in speed and perception of effort than groups who use the radar view first and then the overview.

The hypotheses are tested by looking for effects of View (either main effects or effects in interaction with Task), using three dependent variables (completion time, verbal efficiency, and perceived effort). More detail on the setup of the experiment can be found in (Gutwin and Greenberg 1999).

Findings

A variety of results were obtained, some showing improvement when there was additional awareness information, and some showing no difference between the two displays. We made comparisons of participants using their first interface only (for hypothesis 1), and of participants' differential performance between their first and second interfaces (for hypothesis 2).

For the first hypothesis, where we compared measures from the first interface used, no main effects of View were found, although there were interactions between View and Task. When using the radar view, groups finished the Follow and Direct tasks significantly faster (Figure 11: about 3 minutes with the radar, and about 4.5 with the overview). Also, groups using the radar spoke significantly fewer words in the Follow task (Figure 12: about 100 words with the radar, 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 (Figure 13).

For the second hypothesis, when we compared groups' improvement from their first interface to their second, we found a stronger effect (Figure 14). Using the differential between their first tasks and second tasks, we found main effects of View for completion time differential and for perceived effort differential. That is, and as illustrated in Figure 14, when groups started with the overview and switched to the radar view, their improvement was significantly greater than when they started with the radar and switched to the overview.

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 (Figure 15).

These results add weight to the overall hypothesis that awareness support improves groupware usability. Again, details on the specific results of the study can be found in (Gutwin and Greenberg 1999).

Figure 11. Mean completion times (in minutes)

Figure 12. Mean verbal efficiency (in number of words) for tasks 1-3

Figure 13. Mean questionnaire responses for tasks 1-3

Figure 14. Mean changes in completion time from first to second attempts at a task.

Figure 15. Perceived effort differentials between first and second attempts at a task.

Discussion: how did the awareness information assist performance?

Two underlying issues observed in the experiment warrant further discussion. As outlined in the workspace awareness framework, awareness information is useful for a variety of collaborative actions. Two of these can be seen as primary reasons for the radar view's success in the study: visual awareness information makes it easier to communicate useful information without talking; and awareness information gave people confirmation about the other person's activities.

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. One way that visual information aided the task was by allowing people to use strategies that were better suited to the task and therefore more effective. 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.

The radar view also provided continuous feedback about location and piece position, feedback that 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, where 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 for establishing that instructions had been understood and carried out.

Again, the difference between the overview and the radar was the difference between verbal and visual information. 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 their actions. Clark (1996) summarizes 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 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.

Conclusions

Our main message is: 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. First, designers need a better understanding of what exactly is meant by workspace awareness. This is the role of our workspace awareness framework, where 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 artefacts 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 where 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., and Pew, R., Situation Awareness and the Cognitive Management of Complex Systems, *Human Factors*, 37(1), 85-104, 1995.
- Benford, S., Bowers, J., Fahlen, L., Greenhalgh, C., and Snowdon, D., User Embodiment in Collaborative Virtual Environments, *Proceedings of the Conference on Human Factors in Computing Systems (CHI'95)*, 1995, 242-249.
- Birdwhistell, Ray L., *Introduction to kinesics : an annotation system for analysis of body motion and gesture*. University of Kentucky Press, 1952.
- Borning, A., and Travers, M., Two Approaches to Casual Interaction over Computer and Video Networks, *Proceedings of the Conference on Human Factors in Computing Systems*, New Orleans, LA, 1991, 13-19.
- Brennan, S., *Seeking and Providing Evidence for Mutual Understanding*, Ph.D. thesis, Stanford University, Stanford, CA, 1990.
- Brinck, T., and Gomez, L. M., A collaborative medium for the support of conversational props, *Proceedings of Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'92)*, Toronto, Ontario, 1992, 171-178.
- Clark, H., *Using Language*, Cambridge University Press, Cambridge, 1996.
- Dix, A., Finlay, J., Abowd, G., and Beale, R., *Human-Computer Interaction*, Prentice Hall, 1993.
- Dourish, P., and Bellotti, V., Awareness and Coordination in Shared Workspaces, *Proceedings of the Conference on Computer-Supported Cooperative Work*, Toronto, 1992, 107-114.
- Ellis, C., Gibbs, S., and Rein, G., Groupware: Some Issues and Experiences, *Communications of the ACM*, 34(1), 38-58, 1991.
- Endsley, M., Toward a Theory of Situation Awareness in Dynamic Systems, *Human Factors*, 37(1), 32-64, 1995.
- Fraser, M., Benford, S., Hindmarsh, J., and Heath, C. Supporting awareness and interaction through collaborative virtual interfaces. *Proceeding of ACM UIST99*, 1999, 27-36.
- Gaver, W., Sound Support for Collaboration, *Proceedings of the Second European Conference on Computer Supported Cooperative Work*, 1991, 293-308.
- Gilson, R. D., Introduction to the Special Issue on Situation Awareness, *Human Factors*, 37(1), 3-4, 1995.

- Greenberg, S., and Bohnet, R., GroupSketch: A multi-user sketchpad for geographically-distributed small groups, Proceedings of Proceedings of Graphics Interface '91, Calgary, Alberta, 1991.
- Greenberg, S., Gutwin, C. and Cockburn, A. (1996) Using Distortion-Oriented Displays to Support Workspace Awareness. in People and Computers XI (Proceedings of the HCT'96), A. Sasse, R.J. Cunningham, and R. Winder, Editors. Pages 299-314, Springer-Verlag. Conference held at Imperial College, London, August 20-23.
- Greenberg, S., Gutwin, C., and Roseman, M. (1996). Semantic Telepointers for Groupware. Proceedings of the OzCHI '96 Sixth Australian Conference on Computer-Human Interaction, Hamilton, New Zealand, November 24-27.
- Greenberg, S., and Marwood, D., Real Time Groupware as a Distributed System: Concurrency Control and its Effect on the Interface, Proceedings of the Conference on Computer-Supported Cooperative Work, Chapel Hill NC, 1994, 207-217.
- Gutwin, C., and Greenberg, S. (1999) Effects of Awareness Support on Groupware Usability. *ACM Transactions on CHI*, vol. 6 no. 2, 243-281.
- Gutwin, C. *Workspace Awareness in Real-Time Distributed Groupware*. Ph.D. Dissertation, Department of Computer Science, University of Calgary, 1997.
- Gutwin, C., and Greenberg, S. Descriptive Framework of Workspace Awareness for Real-Time Groupware. Computer Supported Cooperative Work, Kluwer Academic Press, In press.
- Gutwin, C., Roseman, M. and 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, Boston, Mass., November 16-20, ACM Press.
- Heath, C., Jirotko, M., Luff, P., and Hindmarsh, J., Unpacking Collaboration: the Interactional Organisation of Trading in a City Dealing Room, Computer Supported Cooperative Work, 3(2), 147-165, 1995.
- Heath, C., and Luff, P., Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms., Computer-Supported Cooperative Work, 1(1-2), 69-94, 1992.
- Hutchins, E., Cognition in the wild. MIT Press, 1996.
- Hutchins, E., 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, Lawrence Erlbaum, Hillsdale, NJ, 1990.
- Ishii, H., and Kobayashi, M., ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact, Proceedings of the Conference on Human Factors in Computing Systems, Monterey, CA, 1992, 525-532.
- McGrath, J., Groups: Interaction and Performance, Prentice-Hall, Englewood Cliffs NJ, 1984.
- Neisser, U., Cognition and Reality, W.H. Freeman, San Francisco, 1976.
- Norman, D., Things That Make Us Smart, Addison-Wesley, Reading, Mass., 1993.
- Robinson, M., Computer-Supported Cooperative Work: Cases and Concepts, Proceedings of Groupware '91, 1991, 59-75.

- Roseman, M., and Greenberg, S., Building Real-Time Groupware with GroupKit, a Groupware Toolkit, *Transactions on Computer-Human Interaction*, 3(1), 66-106, 1996.
- Salvador, T., Scholtz, J., and Larson, J., The Denver Model for Groupware Design, *SIGCHI Bulletin*, 28(1), 52-58, 1996.
- Seely Brown, J., Collins, A., and Duguid, P., Situated Cognition and the Culture of Learning, *Educational Researcher*(January-February), 32-42, 1989.
- Segal, L., 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, Lawrence Erlbaum, Hillsdale, NJ, 1995.
- Segal, L., Effects of Checklist Interface on Non-Verbal Crew Communications, NASA Ames Research Center, Contractor Report 177639, 1994.
- Short, J., Williams, E., and Christie, B., Communication Modes and Task Performance, in *Readings in Groupware and Computer Supported Cooperative Work: Assisting Human-Human Collaboration*, R. M. Baecker ed., 169-176, Morgan-Kaufmann Publishers, Mountain View, CA, 1976.
- Tang, J., Listing, Drawing, and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams, Ph.D. thesis, Stanford University, Stanford, CA, 1989.
- Tang, J., Findings from Observational Studies of Collaborative Work, *International Journal of Man-Machine Studies*, 34(2), 143-160, 1991.
- Tang, J. C., and Minneman, S. L., Videodraw: A video interface for collaborative drawing, *Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems*, Seattle Washington, 1990, 313-320.
- Tang, J. C., and Minneman, S. L., VideoWhiteboard: Video shadows to support remote collaboration, *Proceedings of ACM SIGCHI Conference on Human Factors in Computing Systems*, New Orleans, 1991, 315-322.
- Tatar, D., Foster, G., and Bobrow, D., Design for Conversation: Lessons from Cognoter, *International Journal of Man-Machine Studies*, 34(2), 185-210, 1991.