Cite as:
Greenberg, S. (1997). Collaborative Interfaces for the Web. In C. Forsythe, E. Grose and J. Ratner (editors), Human Factors and Web Development, Chapter 18, p241-254, LEA Press, ISBN 0-8058-2823-0.

18 Collaborative Interfaces for the Web

Saul Greenberg
University of Calgary

The 1960s heralded a great leap forward in the way people used computers. During that decade, people moved away from batch processing via punched cards and paper tapes to interactive online dialogues. Yet computers were expensive, and the only way to make interactive systems affordable was through *time-sharing*, a method that allowed many people to interact simultaneously with a single computer. A time-sharing system went to great lengths to give each user the illusion that he or she was the only person working on it (Figure 18.1a). Personal accounts and access control lists kept people's individual filing systems separate from each other, all in the name of security. Unfortunately, these measures became barriers to collaboration. Information sharing, although possible, was heavyweight. Permissions had to be set; files had to be transferred explicitly from one user's space to another; information had to be included in mail.

The World Wide Web (WWW)—another great leap forward—has changed this process. Simply by making personal or corporate information accessible to a site's Web server, people around the world can distribute and share information with each other. The WWW and its hypermedia structure give the illusion of one large filing system. However, significant barriers to effective collaboration remain. Although information is a shared resource, standard Web browsers are still single-user tools that partition one person from another and offer little support for people to contact each other and engage in conversation over this information (Figure 18.1b).

In this chapter, I consider how this final barrier to collaboration can be removed from the WWW. I will show how Web browsers can be turned into groupware interfaces that allow people to contact each other, discuss documents, and create artifacts through their displays in real time (Figure 18.1c). Such systems are now being developed, and there is already a proliferation of Internet-based groupware systems that bring together tools including telephony, address books, text-based chat tools, electronic whiteboards, application sharing, and presentation packages. Commercial examples are the Netscape Conference tools for peer-to-peer communications (http://www.netscape.com/comprod/products/communicator), Microsoft's NetMeeting for multi-point conferencing (http://www.microsoft.com/netmeeting/), and Intel's ProShare which includes video conferencing (http://www.intel.com/proshare/conferencing/). While most commercial systems are not yet well-integrated with browsers, we should expect this to happen shortly. The recent availability of Java and powerful plug-in interfaces have made such systems reasonable to implement as part of a Web browser.

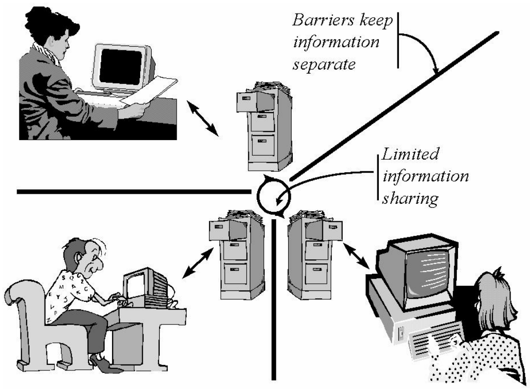

Figure 18.1a. Traditional time-shared computer systems: Barriers keep people and their information separate from each other.

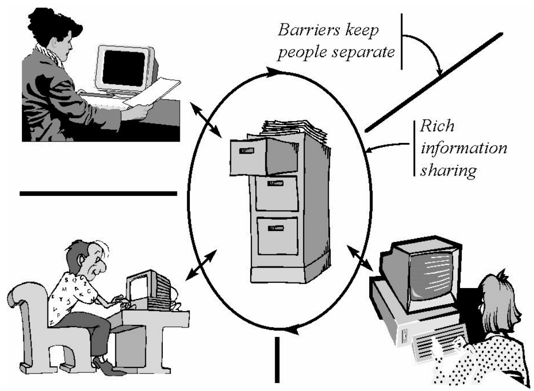

Figure 18.1b. The WWW: Information is easily shared, but barriers keep people from collaborating on it.

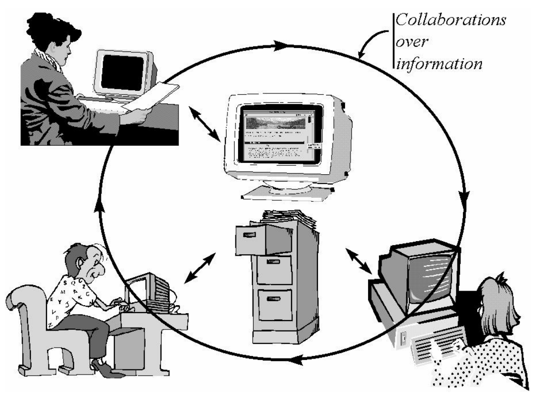

Figure 18.1c. The WWW and Groupware: People can collaborate on shared information in real time.

Figure 18.1. From time-sharing to the WWW to Groupware.

Groupware is not a new idea. Engelbart & English (1968), for example, demonstrated voice and video conferencing as well as screen sharing in the late 1960s. Now, an active, decade-old discipline called Computer Supported Cooperative Work (CSCW) studies how people work together (both with and without groupware), and how computers and related technologies can be designed to support effective collaborations. There are specialized conferences on this topic (ACM Proceedings, 1986–present), edited collections (e.g., Baecker, 1993; Greif, 1988; Greenberg, 1991; Greenberg, Hayne, & Rada, 1995), and even dedicated journals (Schmidt, 1991–present).

Because the scope of CSCW and its literature is too large to summarize in a few pages, the following section provides an introduction and overview of only one area of groupware and its human factors: shared visual workspaces. The subsequent section provides a case study of a shared visual workspace called GroupWeb, a prototype groupware web browser (Greenberg & Roseman, 1996a, 1996b).

Shared Visual Workspaces

In the everyday world, shared physical workspaces (such as whiteboards, control panels, and tabletops) and the artifacts they contain (sketches, controls, documents, structured drawings) act as a stage and offer props for rich person-to-person interaction. Not only are the information artifacts in that space important, but also the ways that interaction over the information is facilitated. For example, people are aware of what others are doing as they are doing it. They can see where others are looking. They can gesture over the surface, can focus other peoples' attention to a part of a workspace; and can guide the presentation of information. They can work individually or collaboratively—even simultaneously—if they want to. Physical workspaces have very specific properties and constraints that afford these things. Creating electronic shared workspaces requires a deep understanding of the affordances of physical workspaces and the way people work over them. This section describes some human factors of a shared workspace.

Studies of Shared Visual Workspaces

Before considering how people interact over a workspace, we first need to know the minimal requirements for effective person-to-person communication. Many studies—for instance, Chapanis, Ochsman, Parrish & Weeks (1972) and Bly (1988)—compared the ways people communicate in real time through technology and contrasted communication conditions such as typed notes, telephone-quality voice, high-fidelity voice, and video links. In all cases, voice communication was found to be important. Although it seems self-evident that people working together in real time over an interactive workspace should have at least a voice channel, many of today's systems provide only a text chat facility, which is inadequate.

A fundamental property of most shared workspaces is the ability to make marks in it. Bly (1988) studied two designers communicating over a drawing surface. From her observations, she asserted that the drawing process—the actions, uses, and interactions on the drawing surface—are as important to the effectiveness of the collaboration as the final artifact produced. She also noticed that allowing designers to share drawing-space activities increases their attention and involvement in the design tasks. When interaction over the drawing surface was reduced, the quality of the collaboration decreased.

Tang (1991) refined Bly's findings through his ethnographic study of small teams solving problems over large sheets of paper. Some of his observations included the following:

- *Orientation:* When people sat around the table, drawings made on the paper were oriented in different directions. Although people had greater difficulty drawing and perceiving the images, orientation provided a resource for facilitating the meeting. Because drawings faced a particular person, a context and an audience were established. Participants' marks that were aligned to an image conveyed support and focus. People working on their own images used orientation as a "privacy" boundary until they were ready to call in the group's attention. The seam when moving between individual and group work was small. (Ishii, Kobayashi, & Grudin (1993) further developed the idea of "seamlessness".)
- *Proximity:* When participants were huddled close to each other around the table, the drawing played a key role in mediating the conversation.
- *Simultaneous access:* With good proximity, a high percentage (45% to 68%) of people's activity around a work surface involved simultaneous access to the space by more than one person.

Tang then built a descriptive framework to help organize the study of work surface activity, where every user activity was categorized according to what action and function it accomplished. The actions accomplished by user activity included:

- *Listing*, which produces alphanumeric notes that are spatially independent of the drawing,
- *Drawing*, which produces graphical objects, typically a sketch with textual annotations attached to it,
- *Gesturing*, which is a purposeful body movement that communicates specific information, such as pointing to an existing drawing.

The functions accomplished by user activity included:

- *Storing information,* which preserves group information in some form for later recall,
- *Expressing ideas,* which involves interactively creating representations of ideas in some tangible form, usually to encourage a group response,

- *Mediating interaction*, which facilitates group collaboration and includes turn taking and focusing attention.

Tang found that the conventional view of work surface activity—storing information by listing and drawing—constituted only around 25% of all functions. Expressing ideas and mediating interaction comprised the additional ~50% and ~25% respectively. Gesturing, which is often overlooked as a work surface activity, played a prominent role, about 35% of all people's actions. For example, participants enacted ideas by using gestures to express them, and gestures were used to signal turn taking and focus the attention of the group.

Another important aspect of collaborative work over a shared workspace is gaze awareness, which is a person's ability to monitor the direction of a partner's gaze (Ishii, Kobayashi, 1992). With gaze awareness, one person can see another person's focus of attention. This includes whom the person looks at, whether they make eye contact, whether one person is attending to what the other is doing, and whether people are tracking each other's actions. Gaze awareness also indicates what objects the person attends to in the shared space, and whether both people are looking at the same thing. Through a series of experiments, Ishii, Kobayashi, & Grudin (1992) noticed that gaze awareness helped people facilitate conversation and gesturing, and helped them move between their interpersonal space (where they were looking at each other) and the shared workspace (where they were looking at objects).

Several researchers continued to investigate how people maintain a sense of awareness about who else is in the workspace, where people are operating, and what they are doing (Dourish & Belloti, 1992; Gutwin, Greenberg, & Roseman, 1996). In a physical workspace, people use peripheral vision, auditory cues, and quick glances to track what goes on around them. These kinds of things help people work together more effectively. Gutwin, Greenberg, & Roseman (1996) identified the notion of *workspace awareness*: the collection of up-to-the moment knowledge a person holds about the state of another's interaction with a workspace. They studied people working together to compose a newspaper layout and noticed the following activities, all afforded by workspace awareness.

- *Mixed focus collaboration*: People shifted their focus back and forth between individual and shared activity, and between different parts of the layout.
- *Lightweight information gathering*: One person gained awareness through rapid glances of another person's working area.
- *Integration of information with previous knowledge*. It was easy to assimilate what a person saw now with what was seen before.
- *Anticipation of another's actions*. A person was aware of what another was going to do next.
- *Using awareness of activity.* A person changed what he or she was doing based on another's activity.

- *Interpreting references*. Sounds, comments, and deictic references of one person were interpreted from another's understanding of where the first was located and what he or she was doing.

Workspace awareness brings another dimension to understanding collaborative interactions. It helps people move between individual and shared activities, provides a context in which to interpret other’s utterances, allows anticipation of others’ actions, and reduces the effort needed to coordinate tasks and resources. But the notion of awareness is a slippery concept. Gutwin, Greenberg, & Roseman (1996) tried to pin it down by describing it as a list of information elements that people may keep track of when they work with others in a shared space (Table 18.1, column 1). These elements can then be related to questions that people ask themselves during group work (column 2).

Electronic virtual workspaces must emulate the affordances of physical workspaces, if they are to support a group’s natural ways of working together. Unlike their physical counterparts, all affordances must be programmed in. This is no easy task. For example, current technology means that some things ordinarily taken for granted are hard to do electronically, such as knowing where others are looking, relating body gestures to items in a workspace, glancing around for awareness, and so on. Consequently, designers must consider and test alternative ways to support what people require; they must understand how the constraints and limitations of technology can be mitigated.

GroupWeb, A Case Study

A prototype groupware Web browser we have developed called GroupWeb (Greenberg & Roseman, 1996a+1996b) illustrates how knowledge of

Table 18.1
Elements of Workspace Awareness

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

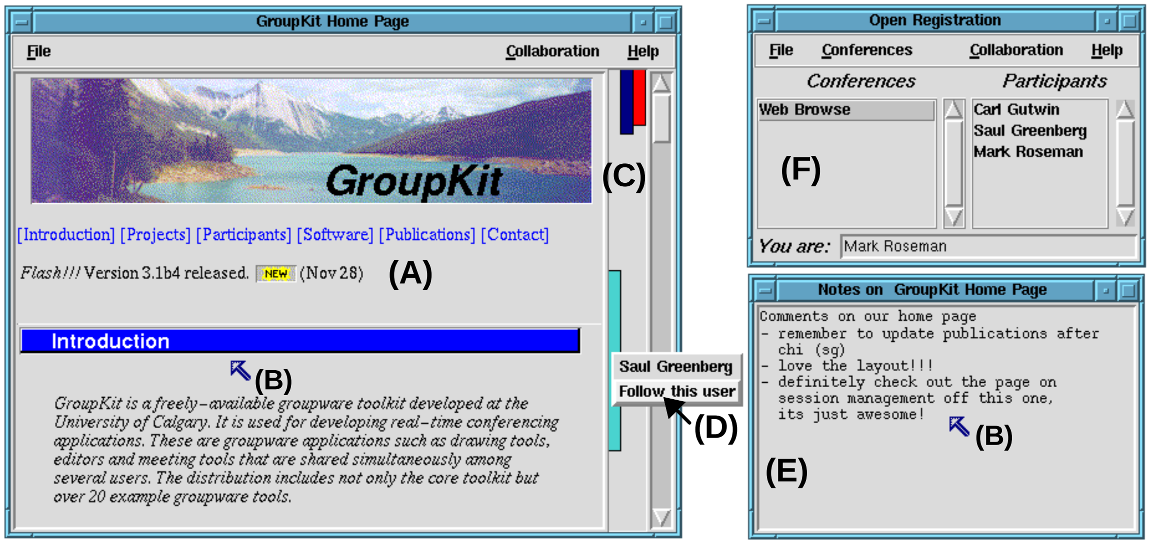

Figure 18.2. GroupWeb, showing (a) a shared web page, (b) telepointers, (c) multi-user scrollbars, (d) view-slaving controls, (e) a group annotation tool, and (f) a session manager for joining the GroupWeb conference.

workspace activities can be applied to the WWW. A first version is illustrated in Figures 18.2 and 18.3, and its design rationale follows. Later, we present a redesigned (but unimplemented) second version that addresses some usability problems we had noticed (Figure 18.4).

Like normal Web browsers, GroupWeb fetches and displays Hypertext Markup Language (HTML) pages. However, several geographically separated people can enter a GroupWeb session via a session manager (Figure 18.2f). Each participant runs his or her own GroupWeb replica, and the main browser window becomes a shared visual workspace (Figure 18.2a). GroupWeb currently has the same limitations as normal Web browsers: Thee pages can be viewed but not altered. Nevertheless, GroupWeb can be used as a tool for remote presentations and for discussing documents over distance.

A basic act of a person using a workspace is selecting material (a drawing or page) and bringing it to the group's attention. In GroupWeb, the material is an HTML page that a group member selects by navigating a link. Whenever any person selects a link to a new page, GroupWeb guarantees that all the browsers in the session navigate to that page automatically by instructing all its replicas to fetch the new page, specified by the Hypertext Transfer Protocol (HTTP) address. GroupWeb does *not* enforce turn taking or any other protocol for selecting links. Rather, it relies on people's natural abilities to mediate interaction (as noticed by Tang, 1991) as well as on the affordances of the workspace (such as seeing other people's actions and their gestures) to make sure that an act is reasonable. Although relying on social instead of technical protocols means that conflicting actions can happen, the gain is the ability of the group to follow its own desires and working styles.

How should the page be presented? The simplest thing would be to keep visuals identical across all displays. However, display sizes and personal desires differ, and it may be onerous to impose this constraint onto distributed meeting participants. GroupWeb "relaxes" strict view sharing by permitting windows to be different sizes, and by reformatting text to fit the display nicely (Figure 18.3). Relaxed view sharing means that all people may not see exactly the same thing (as when orientation differs around a table), but it does provide more flexibility for the way each person wishes to view a page.

A similar question is how scrolling is handled when a page does not fit completely within a window. People sometimes want to pursue individual as well as collective work in a workspace (Gutwin, Greenberg, & Roseman 1996), and they want to be able to go to their own semiprivate space (Tang 1991). GroupWeb allows both independent and synchronized scrolling. With independent scrolling, people can have viewports of different parts of a page, with awareness of others' locations supplied by multi-user scrollbars (Figure 18.2c). Each colored bar represents how much a particular participant can see, as well as the overlap between views, if any. Synchronized scrolling, on the other hand, automatically aligns viewports. To enable this, a menu on each bar is raised (Figure 18.2d), with the menu (and the bar's color) identifying the participant it represents. Selecting 'Follow this user' causes the local display to scroll in synchrony with the other.

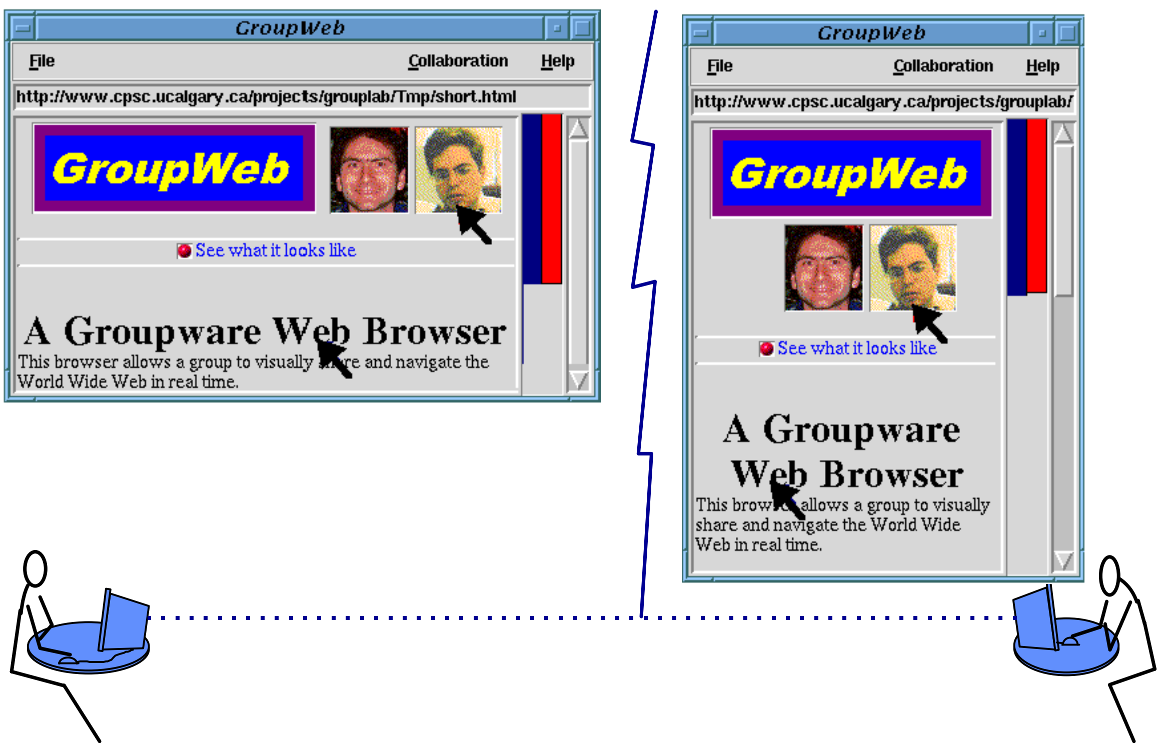

Figure 18.3. Two participant's view of a GroupWeb session, with contents reformatted to fit the different window sizes. Telepointers are on top of the correct text and images, as they are tied to object coordinates, not Cartesian coordinates.

Because synchronization is one way (unless another person also slaves the view), a local user can still scroll quickly to other parts of the page for quick glances.

Gesturing is critical in a shared workspace. As with many groupware workspaces, GroupWeb uses telepointers as a way to transmit and display gestures (Figure 18.2b). All participants have their own telepointers, and they use them to enact ideas, to signal turn taking, to focus the group's attention, and to reference objects on the work surface. Telepointers also afford many awareness elements listed in Table 18.1. For example, their presence and motion indicate those who are present, where they are working, what actions they are doing, and what objects they are manipulating. As people often look at their own telepointers as they are using them, telepointers also offer others a limited form of gaze awareness. Yet implementing telepointers can be problematic when the view of the page can be formatted differently on people's displays. In GroupWeb, we solved this problem by attaching telepointers to letter positions rather than to Cartesian coordinates. Thus the pointer is always over the same text on all displays, as illustrated in Figure 18.3, and most deictic references are correct.

Although the original Web document is not editable, people can attach shared annotations (which Tang (1991) refered to as a listing action) to any page. GroupWeb includes an annotation tool (Figure 18.2e), a multi-user text editor implemented as an add-on workspace. Users can simultaneously enter and edit text at any time, which is displayed on all screens. Characters appear on all displays as they are typed, a process useful for expressing ideas. The annotation is automatically keyed to the current WWW page. Changing to a new page clears the editor; returning to an annotated page restores the text annotation in the editor. Annotations can be used for almost anything: group note-taking, collecting comments, suggesting revisions, and chatting. Telepointers work over this area as well.

GroupWeb still has a long way to go, and using it in practice showed some flaws. In spite of being designed around certain human factors requirements, the way some factors are supported is questionable, and other factors are not catered to at all. We mention this because we constantly have to remind ourselves that translating human factors knowledge into systems is not straightforward—usability testing and iterative development *must* be considered as an integral part of the human factors of groupware design.

1. Although GroupWeb users are expected to be in voice contact, they have to establish this connection as a separate act (e.g., conventional or Internet telephony). Because connecting takes time, it would likely deter short collaborative sessions.
2. A multi-user scrollbar was provided to support awareness of the presence of other people and their location in the document. When we tested the scrollbar in a usability study (Gutwin, Roseman, & Greenberg, 1996), we found that people had problems with these scrollbars simply because it was hard to determine exactly where in the

document the other person was, and exactly what portion of the display was visible to both participants. Although the bars could be used to align views, the information required for fine-grained coordination and awareness was too hard to tease out. Other devices that have passed the usability test could be offered instead. One example is a radar overview that presents a miniature of the entire document and draws rectangles and telepointers on the miniature that correspond to people's individual viewports and telepointers. Optionally, it may just be simpler to enforce strict view sharing by eliminating asynchronous scrolling and by guaranteeing same window sizes.

3. Although text listing is possible through the groupware annotation window, People cannot annotate the page directly by drawing graphics and by attaching text marks to objects. Yet this annotation can form 65% of actions in conventional workspaces (Tang, 1991). GroupWeb should include an annotation layer that sits on top of the shared view. Following techniques applied successfully to groupware sketch pads, people should be able to sketch and type simultaneously on top of a displayed page, should see others' drawing actions as they occur, and should allow annotations to be storable.

Figure 18.4 suggests a simpler (but not yet implemented) redesign of GroupWeb. Although much more constrained than the first version, it should be a reasonable tool that allows people to do remote presentations and discussions over documents. In this version, all windows are kept the same size on all screens, and the pages are formatted identically across all displays. Synchronous scrolling eliminates the need for multi-user scroll bars. Telepointers are still present. An annotation layer has been added so that people can both list and draw on top of the image; this layer eliminates the need for an add-on groupware annotation window.

Summary

This chapter has provided a brief overview of the human factors of shared visual workspaces, along with an example of how shared workspaces can be applied in a simple groupware Web browser. The chapter also includes references for those who require a deeper background in CSCW.

Of course, collaborative interfaces in the WWW go far beyond the shared visual workspaces mentioned here. For example, the chapter does not mention the human factors involved in actually getting a groupware session going. People must be aware of who is around and available for conversation (Kraut, Egido, & Galegher, 1988), must initiate their groupware session, and must bring together their shared material. If any of this is hard to do, then it is unlikely that people will bother to start the session unless there is an overwhelming need (Cockburn & Greenberg, 1993). The design of metaphors that bring people and their groupware

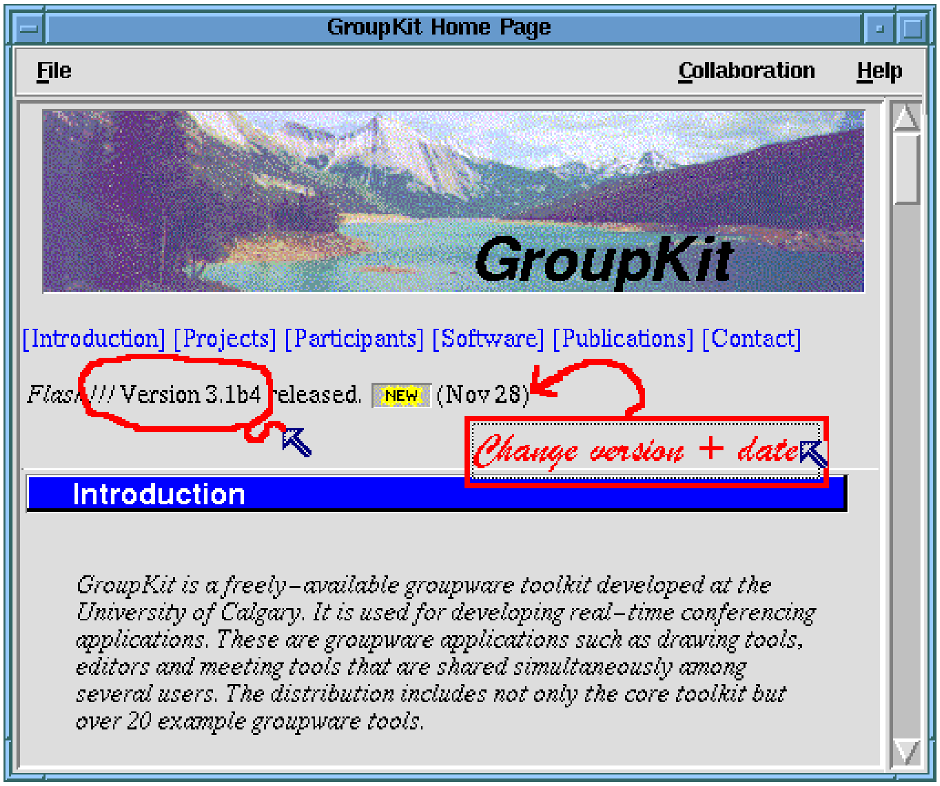

Figure 18.4. A possible simpler but more constrained redesign of GroupWeb. All windows are kept the same size on all screens. Synchronous scrolling eliminates the need for multi-user scroll bars. An added annotation layer eliminates the need for the add-on groupware editor.

together is just as critical as the design of the groupware application. Other issues include security (so that sessions remain private), access control (so that a group can decide who can enter a groupware session), the melding of individual and group work (so that people's individual work can be brought in and out of groupware sessions), the melding of asynchronous and synchronous work, and so on. As well, collaborative interfaces should be designed to support specific tasks. For example, the demands of group authoring have a different (but perhaps overlapping) set of human factors requirements than those found in shared visual workspaces (Rada 1996).

The spreading of the Internet, the popularity of the WWW, the ubiquity of WWW browsers, and the recent availability of Java mean that collaborative interfaces will be built for the WWW. Fortunately, these collaborative interfaces do not have to be crude first efforts with poor usability. We can bring to their design the rich human factors literature and experiences in CSCW and groupware.

References

Baecker, R. (1993). *Readings in groupware and computer-supported cooperative work: Assisting human–human collaboration*, San Mateo, CA: Morgan-Kaufmann.

Bly, S. (1988). A use of drawing surfaces in different collaborative settings. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 250–256), Portland, OR: ACM Press.

Chapanis, A., Ochsman, R., Parrish, R. & Weeks, G. (1972), Studies in interactive communication: I. The effect of four communication modes on the behaviour of teams during cooperative problem solving. *Human Factors,* **14**(6), 487–509.

Cockburn, A. & Greenberg, S. (1993). Making contact: Getting the group communicating with groupware. In *Proceedings of the ACM Conference on Organizational Computing Systems* (pp. 31–41), Milpitas, CA: ACM Press.

Dourish, P. & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 107–114), Toronto, Ont., Canada: ACM.

Engelbart, D. & English, W. (1968). A research center for augmenting human intellect. In *Proceedings of the Fall Joint Computing Conference* (Vol. 33, pp. 395–410), Montvale, NY: AFIPS Press.

Greenberg, S., (Ed.) (1991). *Computer supported cooperative work and groupware*. London, England: Academic Press.

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

Greenberg, S. & Roseman, M. (1996a). GroupWeb: A groupware web browser. In *Video Proceedings of ACM Conference on Computer-Supported Cooperative Work*, Boston: ACM.

Greenberg, S. & Roseman, M. (1996b). GroupWeb: A WWW browser as real time groupware. In *Companion Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 271–272), Vancouver, BC: ACM.

Greif, I. (Ed.) (1988). *Computer-supported cooperative work: A book of readings*, San Mateo, CA: Morgan-Kaufmann.

Gutwin, C., Greenberg, S. & Roseman, M. (1996). Workspace awareness in real-time distributed groupware: Framework, widgets, and evaluation. In R. Sasse, A. Cunningham, & R. Winder (Eds.) *People and Computers XI (Proceedings of HCI'96 BCS Conference on Human Computer Interaction)*, pp. 281-298, Springer-Verlag.

Gutwin, C., Roseman, M., & Greenberg, S. (1996). A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of ACM Conference on Computer Supported Cooperative Work,* Boston: ACM Press.

Ishii, H., & Kobayashi, M. (1992). ClearBoard: A seamless medium for shared drawing and conversation with eye contact. In *ACM Conference on Human Factors in Computing Systems* (pp. 525–532), Monteray, CA: ACM.

Ishii H., Kobayashi M., & Grudin, J. (1993). Integration of interpersonal space and shared workspace: Clearboard design and experiments. *ACM Transactions on Information Systems,* October.

Kraut, R., Egido, C. & Galegher, J. (1988) Patterns of contact and communication in scientific collaboration. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 1–12), Portland, OR: ACM.

Proceedings of the Conference on Computer Supported Cooperative Work. Biannually since 1986, ACM.

Rada, R.(Ed.) (1996). *Groupware and authoring*, London, England: Academic Press.

Schmidt, K. (Editor) (1991 to present) *Computer Supported Cooperative Work: The Journal of Collaborative Computing.* Kluwer.

Tang J. (1991). Findings from observational studies of collaborative work. *International Journal of Man Machine Studies*, **34**(2), pp. 143-160, February.