

# Integrating Communication, Cooperation, and Awareness: The DIVA Virtual Office Environment

*Markus Sohlenkamp*  
GMD (German National Research Center  
for Computer Science)  
Schloß Birlinghoven  
53754 Sankt Augustin, Germany  
Tel: +49-2241-14-2872  
Email: sohlenkamp@gmd.de

*Greg Chwelos*  
Department of Computer Science  
University of Toronto  
Toronto, Ontario, Canada M5S 1A4  
Tel: +416-461-4106  
Email: chwelos@dgp.toronto.edu

## ABSTRACT

DIVA, a novel environment for group work, is presented. This prototype virtual office environment provides support for communication, cooperation, and awareness in both the synchronous and asynchronous modes, smoothly integrated into a simple and intuitive interface which may be viewed as a replacement for the standard graphical user interface desktop. In order to utilize the skills that people have acquired through years of shared work in the real offices, DIVA is modeled after the standard office, abstracting elements of physical offices required to support collaborative work: people, rooms, desks, and documents.

**KEYWORDS:** Groupware, Communication, Cooperation, Awareness, Synchronous/asynchronous, Virtual office, Integration, CSCW

## INTRODUCTION

CSCW and groupware are sufficiently mature that there are now numerous examples of both commercial products and research prototypes for most of the major categories of CSCW technologies (the *unOfficial Yellow Pages of CSCW* lists 340 commercial or experimental groupware systems [26]). For example, there are media spaces for real-time communication [28, 13, 33]; email systems for asynchronous communication; groupware tools for cooperative development of documents, drawings, etc., some synchronous and some asynchronous [2, 3, 16, 23, 30]; software for meeting facilitation [27, 31, 38]; and tools for work flow management [22, 29].

With so many systems available, there is now a trend in CSCW research towards integration along numerous dimensions [1]. One such trend is towards the integration of shared task and person spaces [7, 19]. Another trend is the development of groupware editors which support both synchronous and asynchronous modes of work [2, 4, 14, 15].

Our work continues this integration trend. We have

developed DIVA, a prototype virtual office environment which unites in a single interface many of the key facilities needed by distributed working groups. In particular we have focused on providing integrated support for awareness of co-workers, tools for performing shared tasks, and facilities to support communication, in both the synchronous and asynchronous modes of work. The metaphor of a virtual office provides a consistent interface offering most of the elements needed for group work.

The remainder of this paper is structured as follows. The next section presents the rationale for our work. The main section then describes our prototype, DIVA, followed by a discussion of how well it meets its design goals and a list of opportunities for improvement. Technical aspects of the implementation of DIVA are then briefly presented and the paper concludes with a summary of related work and a discussion of the conclusions that may be drawn from this project.

## DESIGN RATIONALE

We set out to develop a novel multi-user work environment for distributed groups supporting a variety of CSCW activities. Motivating our project are two broad design themes:

- Integrate multiple groupware technologies into a single environment in order to eliminate artificial boundaries between CSCW activities which exist in current products and prototypes but are not found in the real work place. We aim not only to provide multiple facilities, but to smoothly integrate these so that transitions from one to another are almost transparent, allowing users to fluidly move between supporting technologies as their needs vary throughout the day.
- Utilize the intuitions, knowledge, and skills that people have acquired through years of shared work in the real world. We include here making use of users' knowledge of current GUI desktops and their direct manipulation nature.

## Integration

We have aimed for integration along two main, orthogonal dimensions: functionality and time mode. Along the functionality dimension, we targeted three key CSCW services: awareness of co-workers, communication with co-workers, and cooperation on the development of shared

artifacts using groupware tools. The time modes are same-time work (synchronous) and different-time work (asynchronous) [20]. Motivations for these services in terms of the CSCW research agenda have been detailed in other work and therefore will not be reiterated here. However, throughout the remainder of the paper we refer to papers which include motivating evidence and discussion.

	<b>Synchronous</b>	<b>Asynchronous</b>
<b>Communication</b>	Communicate in real time.	Leave messages for others.
<b>Cooperation</b>	Simultaneous work using groupware tools.	Turn-taking work.
<b>Awareness</b>	What are others doing now?	What have others done recently?

Table 1: Six classes of CSCW activity.

These two dimensions interact to form a conceptual framework with six cells, as shown in Table 1. Each cell constitutes a unique class of activity found in group work. Table 2 expands these activity classes, listing specific design goals for each class, plus others which span the classes. The lists are not exhaustive, but rather focus on key issues.

Finally, we aimed for a seamless integration of individual and group work so that the same artifacts, tools, and actions should be used in both cases. For example, the action required to open a document for editing should be the same regardless of the number of users involved.

### Utilize Existing Skills

Our second design theme relates to the look and feel of the interface, rather than the functionality. It guides how the design goals are realized. In order to make use of existing user skills, we wanted the elements of the interface and the operations performed on them to have real world counterparts wherever possible. However, it was clear that this design goal could not be followed rigidly, otherwise some limitations of the real world would be re-created in the virtual office, such as the limitation that physical objects can only exist in one place at a time. Instead, real world elements (objects and actions) are mimicked whenever it is sensible to do so and novel virtual world elements are defined as needed.

As in the single user desktop metaphor, DIVA is a very simple abstraction which models only the essential elements of the real world in the virtual world: people, rooms, desks, and documents. In DIVA, CSCW activities become an integral part of the overall environment. Group actions are not initiated by ad hoc actions in special applications (e.g. make connection in a media space application or start a group session in a shared editor), but arise naturally out of ordinary actions in the shared virtual

	<b>Synchronous</b>	<b>Asynchronous</b>
<b>Communication</b>	<ul style="list-style-type: none"> <li>• make and break verbal and visual contact with one or more other people</li> <li>• form and dissolve private conversations while in a larger conference</li> <li>• quickly suspend all audio</li> <li>• control access to communication spaces</li> </ul>	<ul style="list-style-type: none"> <li>• leave notes for others wherever they are needed</li> </ul>
<b>Cooperation</b>	<ul style="list-style-type: none"> <li>• manipulate (create, edit, etc.) shared artifacts</li> <li>• create, join, or leave synchronous cooperation sessions</li> <li>• select the appropriate degree of coupling for synchronous sessions</li> </ul>	<ul style="list-style-type: none"> <li>• coordinate turn-taking work</li> <li>• determine what changes were made to a shared artifact by others</li> <li>• merge divergent work on shared artifacts</li> </ul>
<b>Awareness</b>	<ul style="list-style-type: none"> <li>• obtain some idea of what co-workers are doing</li> <li>• ascertain a co-worker's availability for contact</li> <li>• control their own level of availability</li> <li>• control the information about themselves which is broadcast to others</li> <li>• know when shared documents are in use by others</li> <li>• know exactly what others are doing during a shared editing session.</li> </ul>	<ul style="list-style-type: none"> <li>• determine when shared artifacts have been changed by others</li> <li>• determine how those artifacts have changed</li> <li>• determine when and where others have left messages for them</li> </ul>
<b>Others</b>	<ul style="list-style-type: none"> <li>• quickly join others in a shared work context (communication and collaboration activities)</li> <li>• suspend and later resume a complete work context</li> <li>• control access to shared artifacts</li> <li>• dynamically configure work groups to adapt to changing needs</li> <li>• customize their view of the shared environment</li> </ul>	

Table 2: Design goals for each class of activity plus others which span classes. Each goal completes the sentence: *Users should be able to \_\_\_\_\_*.

office (e.g. communication channels open when users enter the same virtual room and synchronous group sessions occur when users simultaneously work on the same document).

Finally, in keeping with accepted principles of interface design, we tried to make all operations be direct manipulations of visible artifacts [34]. Our overall approach is a blend of analogical design and evolutionary design [32].

The following section describes the prototype we developed based on this design rationale.

## THE DIVA VIRTUAL OFFICE

The standard office was used as the organizing metaphor for our prototype because it is very familiar to most users and provides the elements necessary to support group work. We do not model the facilities of a real office in a highly realistic way, but rather abstract those elements which best support our design goals. The result, the DIVA virtual office environment, may be viewed as a prototype replacement for the ubiquitous graphical user interface (GUI) desktop. Whereas the GUI desktop is a work environment for a single computer user, DIVA is an environment for group work. DIVA includes many of the features of the single user desktop, but defines a more expansive metaphor, that of the office building, in order to meet the much broader needs of groups.

### Elements of the Virtual Office

To keep the interface simple, DIVA models only four principal elements of real offices: people, rooms, desks, and documents. Additional minor elements are briefcases, stick-on notes, and a trash can.

*People* represent the users of the DIVA system and are implemented as small snapshots with a name beneath. They can be moved by their owner to change their virtual location and thereby determine what they are doing and whom they are in audio/visual contact with.

*Documents* represent the artifacts people work on in the virtual office. They are very similar to documents in the single user GUI desktop, but have some additional properties, such as change status and current activity status, which are needed to adequately support group work. They differ from real documents in that they may be in two or more places at once, for the sake of convenience. The multiple copies in different places are simply access points for a single shared artifact.

*Desks* serve a variety of purposes: they are the arena for work within a room, they control the coupling mode of a cooperation, and they can be used to preserve working context.

*Rooms* are containers for people, desks, and documents. They also control the audio/video communication status of users. Just as people located in the same real room are able to see and hear one another, so too can people in the same DIVA virtual room hear and see each other; when a DIVA user enters a virtual room already occupied by one or more users, audio and video links are established between the

newcomer and the other occupants. Rooms also serve to indicate availability and communication willingness: they can be in different states, providing different levels of access and visibility of their inhabitants. Rooms can be used as private offices, public meeting places, or special purpose spaces.

Rooms themselves are contained in the DIVA virtual office environment. Users may customize their virtual office by selecting their set of potential cooperation partners and placing the rooms as they like. A glance at the rooms contained in the virtual office shows users who is inside each open room.

People move through the virtual office, from room to room, by dragging their icon. They may glance into rooms which are not open and they enter a room in order to do work or establish contact with others inside. Once in a room, the desks and documents in it become visible. By moving shared documents and themselves to a desk in the room, users may work together in either tightly coupled or loosely coupled mode. People working at the same desk are engaged in a focused collaboration, using a tightly coupled editing mode, while people working on the same document but at different desks are less focused and use a loosely coupled mode.

### A DIVA Session

A typical DIVA session is illustrated in Figure 1 (henceforth referred to as ‘the example’). The virtual office, shown from the point of view of user Markus, is displayed in two main windows. The first window (in the background) contains the virtual office itself and the second (in the mid ground) shows the virtual room that the user is currently in. The *virtual office window* displays the DIVA rooms and the people present in the virtual office. It indicates the locations of the users and displays their movements from room to room as animations. Users organize the DIVA rooms as they like by simply dragging the rooms with the mouse.

As illustrated in the example, a glance at the virtual office window provides a broad level of awareness of co-worker activities: Markus, Cici, and Mike are together in Markus’ office; Claus and Andreas have met in the project room; Greg and Thomas are each alone, but available for contact; the people in the “Conference” room would like some privacy; and user Jim does not want to be disturbed, as indicated by the lock on his DIVA office. Other activity cues displayed in this window are described in the next section.

The other DIVA window, labeled ‘Room Markus’ in the example, is the *virtual room window*. It reveals the contents of the room that the user is currently in. In addition to the people who are in the room, the desks and documents in it are shown. Awareness information is also conveyed by this window: Markus, Cici, and Mike are all working on the shared drawing “figure” (shown in the Graphic Editor window); Mike is also editing a spreadsheet named “Costs”; the documents “Song” and “text” have changed since Markus looked at them; and there are three notes for Markus in the room. Working at the same desk,

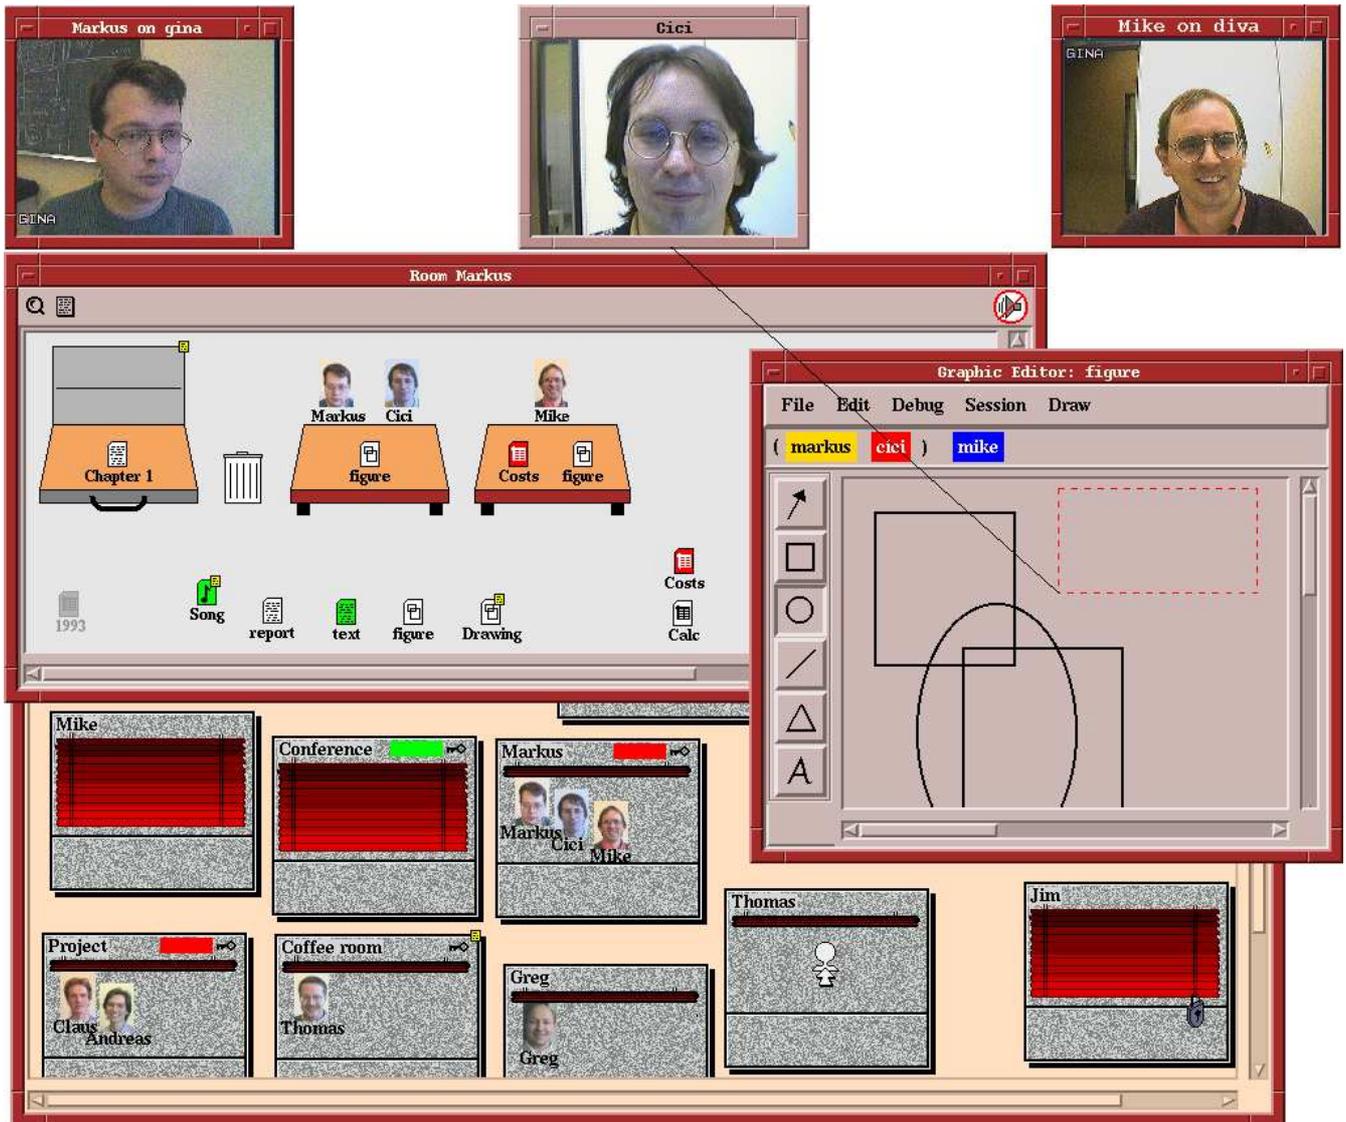


Figure 1: The DIVA Virtual Office. The large background window is the virtual office window, containing rooms and people. The large mid ground window is the virtual room window, showing the contents of a room: people, desks, and documents. The foreground window is a shared graphic editor, currently in use by the three people shown in the video windows at the top. Cici, working at another workstation, is in the process of adding a rectangle to the drawing.

such as Markus and Cici are doing, indicates focused work and results in a tightly coupled synchronous editing session. Although Mike is also editing the same drawing, he is at another desk so his work is independent and he is loosely coupled with the others. In this manner work on the shared artifacts in the room focuses around the desks in the room, as it usually does in real offices.

Finally, the small windows shown in the example are live video images of the people in the room.

## INTERFACE DETAILS

### Synchronous Communication

Real-time person to person communication is supported in DIVA through audio/video conferencing. Control of conferences is based on a very simple model taken from the

real world—people in the same room can see and hear one another while others cannot. So, in order to converse with another person in the DIVA virtual office users simply drag their icon into the DIVA room where the target person is working, using the virtual office window. Small video windows are then automatically opened and placed at the top of the screen, one window for each occupant of the room. Audio channels are opened at the same time and people already in the room are informed of the arrival of a newcomer by an audio cue. The audio and video channels to a particular user are broken when that user leaves the room. In this manner DIVA maintains separate audio/video conferences for each room in the virtual office.

Amongst groups of people holding a general discussion it is often desirable to be able to have private conversations, either to make aside comments or for more extended

conversation. Users at the same desk may initiate a private conversation by dragging their icon so that it overlaps the icon of another user. This action is intended to be an intuitive surrogate for the real world action, in which people talking privately draw close together. It is performed inside a room, in the virtual room window, and is only visible to occupants of that room. During a private conversation in DIVA the sound of the conversation is transmitted to others in the room at a very low volume, while the sounds from the others are received normally. This mimics what happens in real groups.

At times it is necessary to temporarily interrupt video conferencing activity, in which case a user will want to quickly suspend both incoming and outgoing audio. Such interruptions may be caused by a telephone call or the entrance of a real person into a user's physical office. In DIVA audio is temporarily suspended by clicking on the privacy button (the icon on the right end of the tool bar in the room window). Visual feedback indicating this state is provided to all occupants of the room by replacing the user's video image with a static drawing of a person on a telephone.

A possible problem with the virtual office metaphor used in DIVA is that it could be difficult to find a virtual co-worker, if they are in an unusual place or in a closed room. To eliminate the need for exhaustive searches, DIVA provides each user with an own office, called a *home room*, and provides a *find* function. To locate a person, users drag their icons into that person's home office and onto the find icon (the icon shown in Thomas' room in the example). This will move them in the virtual office space to the threshold of the room where the target person is located. Home offices themselves should always be easy to find as users organize the room layout as they like. Users who do not want to be found can disable this feature by locking their home rooms.

### **Asynchronous Communication**

DIVA support for asynchronous communication is based on another object from the real world office: the stick-on note. Notes can be attached to the objects in the virtual office: people, rooms, desks, and documents. To do so, a user drags the note tool onto the target object (the second icon on the tool bar of the room window in the example). This causes the note editor to pop up which both permits creating new notes to attach to the object or reviewing existing notes on the object. After the new note is created, a note icon appears on the object, with one exception. Just as we do not actually stick notes on people in the real world, in DIVA notes directed at people do not appear on their icons but instead appear on their briefcases, where they are both private and accessible to the recipient. Notes can be read by clicking on them or by using the note tool. The note editor and viewer supports text, audio, and video notes.

### **Groupware Applications**

The physical office is again used as the model for the DIVA implementation of shared work; work happens at desks. To edit a document found in a virtual room, the DIVA user drags the corresponding icon to a desk. This causes the

application to start and open the document. This is just like double-clicking on an icon in the typical single user desktop and, for the sake of backward compatibility, double-click has the same effect in DIVA. Documents are closed by double-clicking on their icon on the desk or by closing them from within the application.

*Synchronous* Desks do more than simply start applications, however; they are used as a natural means to control the coupling of synchronous groupware sessions and to save work contexts. To work together in tightly coupled mode, two DIVA users simply work on the document at the same desk. To work together in loosely coupled mode they work on the same document but at different desks—they each drag the document to a different desk. They can easily switch coupling modes by moving from desk to desk.

In the example Markus and Cici are tightly coupled because they are editing the drawing at the same desk, but Mike is loosely coupled because he is working on it at another desk. The number and status of people synchronously editing a document is also indicated within the DIVA groupware applications, on the status line found near the top of the application window. This shows all users, their coupling status (tightly coupled groups are shown in brackets), and their signature colours. Users can change the coupling mode from within the application by dragging their names on this status line. This action also causes their icon to move in the DIVA room window, so that the DIVA room always represents the actual user activity.

Desks can also be used to preserve a working context simply by leaving documents on them. When users drag their icons away from a desk with documents on it, the application windows for the documents close, as if the documents themselves had been closed. The users may then return to working on those documents by later dragging their icon back to the desk, at which point all of the documents on the desk will be reopened.

Unlike most other groupware systems, no special actions are necessary to start or join a synchronous groupware session. Furthermore, the same action—dragging a document to a desk—is used for asynchronous work. The model employed to distinguish work in the two modes is based on physical shared artifacts. If two or more users wish to use the same artifact at the same time, then they must share it. Otherwise users manipulate the artifact on their own, asynchronously. Therefore whenever a user opens a document in DIVA which is already in use by another person, a synchronous session is started. The transition from asynchronous to synchronous is automatic, requiring no other user action. Other people also wanting to work on the document automatically join the synchronous session.

DIVA goes beyond the office metaphor by permitting multiple symbolic copies of documents. This means that a document may be in two or more places at once and may be open in more than one place. The automatic creation of groupware sessions among users simultaneously editing a document occurs even if those users are in different virtual rooms accessing different symbolic copies of the document.

In this case the users have no accompanying audio/video conference and they are necessarily loosely coupled because they are at different desks. This implementation of symbolic copies allows users to keep documents in as many places as is convenient. For example if A and B are writing a paper together, they may each have a copy in their office and they may also have a copy in their project room, where they often work together. This automatic creation of synchronous editing sessions during simultaneous access not only smoothly integrates the two time modes of work but it also minimizes the occurrence of parallel, divergent work.

In summary, the DIVA virtual desk, as the focus of work, supports: synchronous, asynchronous, and individual work; easy switching of coupling modes; and dynamic extension and contraction of synchronous sessions.

Closely related to the desks, every user has a personal briefcase which is always present in the virtual room window (at the upper left of the room window in the example). Unlike the other contents of the window, which change according to the room the user is in, the briefcase does not—it moves from room to room with the user and is not visible to others, though they see their own. The briefcase may be used to move documents from room to room, as a place to privately edit documents while in a room with others, and as a consistent personal working context which moves from room to room.

The actual tools for working in DIVA are multi-user applications built with our GINA application framework [4, 37]. A wide variety of prototype multi-user applications have been implemented in GINA in order to demonstrate its generic nature. These applications include a text editor, spreadsheet, structured drawing tool, music editor, and a chess program. Facilities to support synchronous group editing which are provided by every multi-user GINA application include group awareness in the form of visual representation of others' actions, unlimited multi-user undo/redo, multiple coupling modes, embedded annotations, optimistic concurrency control, and conflict resolution.

**Asynchronous** Asynchronous cooperation is also supported in DIVA, particularly with catch-up and merge facilities. A key need of people editing documents asynchronously is the ability to know what changes others have made in their absence. *Catch-up* facilities which provide this sort of awareness have been developed using a variety of approaches: manual embedded annotations (e.g. MS-Word), automatic annotations such as change bars or red-lining (e.g. MS-Word), versioning with automatic differencing [e.g. 30], mark-up systems [16], and history review systems [39]. With the exception of [39], which focuses on shared drawing, this work has addressed only the text editing case.

DIVA, on the other hand, provides a uniform mechanism for catch-up across all multi-user GINA applications, based on the replay of saved history. Changes made by others are replayed with animation so that they may be viewed exactly as if the user had been there watching them being made, except that the replay may be sped up. This mechanism is

provided by the multi-user GINA framework itself, so that all applications developed using GINA inherit the catch-up facility, including our basic spreadsheet, text editor, and graphics editor. In fact, the exact same mechanism which reveals the actions of others during synchronous editing is used to show those actions during asynchronous catch-up.

DIVA users know from the visual cues on the document icon if a document has changed before they open it. On opening a changed document, users are given the opportunity to catch-up to the changes before proceeding with their editing. Furthermore, the same replay mechanism can be used at any later time to review changes made by others or by oneself.

Merging of parallel, divergent work is also assisted by the GINA merge facility which semi-automatically appends one stream of changes after another. As part of our on-going work we are exploring advanced facilities for both the catch-up and merge operations.

### **Awareness**

Several recent research projects have focused on providing general awareness within a distributed work group in order to facilitate collaboration [6, 9, 10, 11, 13]. DIVA too bolsters group awareness through a variety of mechanisms at different levels of detail and in both synchronous (What are others doing?) and asynchronous (What did others do?) modes.

**Synchronous** The virtual office window provides a broad overview of co-workers' activities throughout the virtual office, while the virtual room window provides more detailed information about a particular room. In the virtual office, it is possible to observe the location of other users, their movement, and who is talking to whom, as well as determine the availability of others for contact. Much of the information may be perceived even when not being actively attended to, such as the animated movement of people and documents. Additionally, rooms containing documents which the user has access to and which are currently in use are flagged with a red bar at the top of the room (the "Project" room in the example).

In a virtual room, users can see who is working with whom and on which documents. The coupling mode between users is also indicated clearly, providing an important clue as to how focused a work session is and thus how available the participants are for interruption. The document icons visually indicate the status of the document: documents not opened by the user but which others are editing, even in a different room, are flagged as in-use and displayed with a red background (e.g. "Costs" in the example). Additionally, when sharing a room, the audio and video channels provide further direct cues as to the activities of co-workers. DIVA also provides artificial audio cues for certain actions, such as another user entering the room or joining a groupware session.

Finally, while editing a shared document using a multi-user GINA application, the activities of others can be seen, especially in tightly coupled mode. The status line of these applications shows the names of all the connected users (in

unique, signature colours), together with their coupling mode. To a varying degree, depending on the coupling mode, the operations performed by each user are animated for others in that user's colour. Furthermore, temporary rubber band lines drawn from the video window of the user issuing a command to his cursor position provide an unobtrusive visual link between a command and its originator. In the example Cici is currently drawing a rectangle which Markus can see grow in real-time as Cici drags his mouse. The outline is displayed in Cici's signature colour and the cursor point is connected to his image by the temporary rubber band line.

All of these cues—in the virtual office, in a virtual room, and within a shared application—provide information as to what co-workers are currently doing.

**Asynchronous** DIVA also provides information as to what others have done. In the virtual office window, rooms containing documents which the user has access to and that have been changed by others are flagged with a green bar at the top of the room (the "Conference" room in the example). This flag is also set if notes have been left for the user on artifacts in the room. In the virtual room window, documents which have been changed by others since the user last edited them are flagged as changed and displayed with a green background (documents "Song" and "text" in the example). Notes left on objects by others are shown as yellow squares on the corner of the objects. Notes are on the "Coffee room", the briefcase, and the "Song" and "Drawing" documents in the example. Visual cues are thus provided both at the office level and at the room level indicating what others have done that is of interest to the user.

In summary, DIVA supports awareness of co-workers activities, both synchronous and asynchronous, at a variety of levels.

**Privacy** We believe that the DIVA awareness information, excluding the audio and video channels, is benign and unthreatening and thus is unlikely to raise concerns over privacy. Nonetheless it is possible for users to limit this information by closing or locking the room they are currently in. This does not overly restrict their ability to work as they can always bring copies of documents into the room with them. Users can also use the access control mechanisms to limit access to rooms and documents. As usual, the cost of privacy is loss of awareness and accessibility.

Working in a DIVA room, on the other hand, necessarily entails a significant loss of privacy as video and audio channels are automatically established. However, even these can be moderated, through the private conversation mechanism and the privacy facility which temporarily stops all DIVA audio activity and outgoing video. Careful study of the behavior of groups using DIVA over a period of time is required to determine if our hunches are correct or if further effort needs to be expended in order to find a set of awareness mechanisms which better afford the balancing of accessibility and privacy.

## Access Control

Access control, in one form or another, is an essential part of any practical multi-user environment. DIVA provides availability states for rooms which give users control over both their availability and the awareness information about them which is conveyed to others. It also includes access lists to give users control over the use of rooms and documents.

**Availability Status** In a manner similar to the door states use in the Ontario Telepresence Project [8], DIVA allows users of rooms to signal and to limit their availability for contact. Rooms may be *open*, *locked*, or *shuttered*. Open rooms may be entered by anyone and their occupants are visible in the virtual office window. This state signals high availability for contact and provides the highest level of awareness. Locked rooms can only be entered by those with a key to the room (see below) and the occupants of the room cannot be seen without entering the room. The locked state indicates very low availability for contact and provides the maximum degree of privacy, at the expense of awareness. The shuttered state provides an intermediate state between these extremes. The occupants of a shuttered room can not be seen directly but can be seen by moving to the threshold of the room. This causes the blinds to lift momentarily, sends an audio cue to the room occupants, and allows them to see in their DIVA room windows who is glancing in. These access states are indicated visually, as illustrated in the example; most of the rooms are open, but Jim's room is locked and both Mike's room and the "Conference" room are shuttered.

**Access Lists** DIVA implements rudimentary object access control in the form of *access lists* for rooms and documents. The room access list determines which users are allowed to enter the room when it is locked. Conceptually, users on the access list have a key to the room. Visual feedback is provided in the form of small key icons on the rooms which the user has access to (e.g. the "Coffee room" in the example). Only users on the access list for a room are able to change the room access status. Typically, the only person on the access list of a private office is the owner of that room.

The document access list controls document appearance and accessibility. Users on the access list of a document will see the document icon in its normal form in the virtual room window and have full access to it. Users not on the list see a mute gray equivalent of the document icon and have no access to the document (document "1993" in the example).

Room and document access lists are managed similarly. Anyone on an access list may add others to the list while initially the list only contains the person who created the corresponding room or document. Finally, anyone on the access list of an object may set a special flag granting universal access to the object. In this case all users may access the object. While much more complex access control mechanisms are possible, this simple mechanism is all that is needed in our prototype. It allows both the evolution of social protocols for access control similar to those used in real offices and the exploration of related interface issues.

## Object Properties

The room availability status and the object access lists are properties of the objects in the DIVA interface. Rooms and documents each have other properties as well. These can be viewed and changed with the *property tool* which is the magnifying glass at the left end of the tool bar in the example. Dragging this tool onto an object pops up the property dialog box for that object from which the properties can be viewed or changed.

## EVALUATING THE DESIGN

Table 3 lists services provided by DIVA for each of the six CSCW activity classes defined in Table 1, as well as other services which cut across these classes. To some degree, each of the design goals listed in Table 2 have been met. However, there are numerous opportunities for

improvement since most of the goals are open-ended in nature. Some of these opportunities, identified as limitations of our current implementation, are listed in Table 4.

Regarding our second design theme, the utilization of existing user skills through analogical design, most of the DIVA elements have clear real world counterparts. These include the interface objects (people, rooms, desks, documents, briefcases, and notes) and many interface operations: entering and leaving audio/video conferences, forming private conversations, opening documents, entering and leaving groupware sessions, controlling coupling in groupware sessions, determining who is where and with whom, controlling availability for contact, leaving messages, access control on rooms. Other elements which are embedded in the context of the metaphor but are

	Synchronous	Asynchronous
<b>Communication</b>	<ul style="list-style-type: none"> <li>• Provides flexible audio/video conferencing.</li> <li>• Controlled by simple operations based on office metaphor.</li> <li>• Supports private conversations.</li> <li>• Supports complete suspension of audio.</li> <li>• Provides control of access to conferences (rooms).</li> </ul>	<ul style="list-style-type: none"> <li>• Supports messages left: a) within a shared artifact, b) on an office object, and c) for the user.</li> </ul>
<b>Cooperation</b>	<ul style="list-style-type: none"> <li>• Supports flexible shared editing sessions.</li> <li>• Provides awareness of others' edit activities; who is doing what.</li> <li>• Supports different coupling modes and makes coupling visible.</li> <li>• Provides unlimited undo/redo.</li> <li>• Provides concurrency control to manage conflicting simultaneous changes to artifacts.</li> </ul>	<ul style="list-style-type: none"> <li>• Transparently coordinates access to shared artifacts.</li> <li>• Provides a pervasive messaging facility to coordinate asynchronous work.</li> <li>• Provides catch-up via activity reply.</li> <li>• Provides support for merging divergent work by merging activity streams.</li> </ul>
<b>Awareness</b>	<ul style="list-style-type: none"> <li>• Shows virtual location, grouping, and availability of co-workers (office view).</li> <li>• Shows who is editing what, including the degree of focus of connected groups (room view).</li> <li>• Permits users to control this information.</li> <li>• Provides visual and aural awareness, and immediate access, similar to that found in open offices (open audio/video connection).</li> <li>• Provides explicit audio cues to signal consequential activities, such as entering a room.</li> <li>• Shows current activity, such as documents currently in use and rooms containing these.</li> <li>• Provides detailed feedback of others' activities during synchronous editing sessions.</li> </ul>	<ul style="list-style-type: none"> <li>• Signals recent activities of others, such as documents changed by others and rooms containing these.</li> <li>• Signals messages left by others: a) within a shared artifact, b) on an office object, and c) for the user.</li> <li>• Shows changes made to shared artifacts by others through activity replay.</li> </ul>
<b>Others</b>	<ul style="list-style-type: none"> <li>• Provides a means to save a work context which can later be restored (desks).</li> <li>• Provides shortcuts to movement between work contexts (rooms).</li> <li>• Provides shared artifacts with access control and multiple access points (symbolic copies).</li> <li>• Transparently coordinates access to shared artifacts.</li> <li>• Provides private work spaces (briefcase and home room) in parallel with the shared work spaces.</li> <li>• Supports changing group needs through the simple creation and modification of work spaces (rooms).</li> <li>• Allows individuals to tailor their view of the shared environment (office view).</li> </ul>	

Table 3: Services provided by DIVA for each of the CSCW activity classes plus other services which span these classes.

somewhat ad hoc are: the notification of document state, access control on documents, symbolic copies of documents, and object properties. Finally, two DIVA elements which are outside of the metaphor altogether are the privacy operation to suspend audio and the facility to find another person.

In summary, despite the limitations listed in Table 4, DIVA does successfully manifest our two main design themes. First, it smoothly integrates communication, cooperation, and awareness in both the synchronous and asynchronous modes. Second, it does so almost entirely within the unifying model of the virtual office so that the elements and operations of the interface are easily understood.

Casual observation of people using an early version of DIVA during an open house revealed that they were able to make use of it with only minimal help from one of the authors who was sitting in another room and interacting with the users via DIVA. Although much more formal evaluation is needed, it would appear that users do in fact find the interface simple and intuitive.

**IMPLEMENTATION**

The virtual office and the multi-user applications that it supports are based on the application framework GINA [36]. Implemented in Common Lisp and CLOS (Common Lisp Object System), GINA encapsulates each user action as a *command object*. Early versions of GINA were targeted at single user applications, but the command object approach made it easy to extend GINA to provide a generic framework for multi-user applications with a replicated architecture [4]. Cooperating copies of a multi-user GINA application communicate by exchanging command objects, so the same command which is executed locally is also executed remotely by all participating copies of the

application. Optimistic concurrency control is used in order to provide very fast local feedback. Conflicting user actions are detected by the use of virtual time stamps and are resolved by reordering the conflicting commands using the undo/redo mechanism.

Multiple coupling modes are provided within GINA by constraining the types of commands that are broadcast. For strict WYSIWIS (What You See Is What I See) work (tightly coupled), all commands, including selections, the scrolling of windows, etc., are transmitted, thus forcing all users to have the same view of the document. For loosely coupled work, where only the document state is to be shared, just state-changing commands are transmitted. Users may have differing views of the document because the view-changing commands are not transmitted to others.

While the multi-user GINA applications are based on a replicated architecture (each user runs their own copy of the application), DIVA itself is currently centralized. All data about the virtual world is contained in a single database, and all DIVA applications are started from the same Lisp process. In the future, the implementation will be changed to match the replicated model provided by the GINA framework.

Unlike most media spaces which use separate analog networks for sending audio and video [e.g. 13, 28, 33], DIVA transmits these signals over the same TCP/IP network used for computer communications. Audio connections are provided by using AudioFile audio servers [24] and special client applications. The servers support the mixing of multiple input channels, which permits DIVA to combine voices from other users with its own audio cues. While the audio quality is adequate for our prototype, it is not good enough for regular use. It is poor because TCP/IP was simply not designed to carry continuous time-critical

	<b>Synchronous</b>	<b>Asynchronous</b>
<b>Communication</b>	<ul style="list-style-type: none"> <li>• Has the usual limitations of video conferencing—paucity of non-verbal cues to coordinate turn-taking in conversation and very little information about work context is conveyed [17].</li> <li>• Has the usual limitations of non-spatial audio—it is difficult to determine who is speaking or to sort out simultaneous speech.</li> <li>• Audio quality is poor and video bandwidth is limited due to current network technologies.</li> </ul>	<ul style="list-style-type: none"> <li>• Uniform status of notes—no ability to signal varying importance or urgency.</li> </ul>
<b>Cooperation</b>	<ul style="list-style-type: none"> <li>• The supported applications are toy prototypes—does not work with existing tools.</li> </ul>	<ul style="list-style-type: none"> <li>• Has no facilities for version control.</li> <li>• Catch-up and merge support is generic, but rudimentary.</li> </ul>
<b>Awareness</b>	<ul style="list-style-type: none"> <li>• Although DIVA provides users with awareness of activities in the DIVA virtual world, it does not adequately reveal activities occurring in the real world outside of the DIVA environment.</li> </ul>	<ul style="list-style-type: none"> <li>• Although DIVA signals that documents have been changed by others, at present it provides no summary information as to the author, type, and magnitude of those changes—one must perform the catch-up process to learn this.</li> </ul>
<b>Others</b>	<ul style="list-style-type: none"> <li>• Extensions are required to scale DIVA up to a very large user group.</li> </ul>	

Table 4: Some limitations of the implementation.

data streams. The advent of ATM networks will remedy this limitation.

For the video conferences, a miniature camera attached to the top of each workstation sends an analog video signal to a special video board which displays the image in an X window. A video server process sends the contents of this window to the video servers on the workstations of the other members of the video conference. The video frames are compressed by the video board using JPEG compression prior to transmission. The video servers run at a lower priority so that when the network becomes loaded, the video frame rate drops before the audio or groupware performance degrades.

### **RELATED WORK**

This version of DIVA evolved from a much simpler first implementation [35]. Several other systems bear similarities to aspects of DIVA. The VOODOO system [25], a prototype virtual open office environment, provides persistent audio and video connections such as exist in DIVA virtual rooms. The Xerox Rooms system [18] uses virtual rooms to organize the work of individual users. The DIVE system [12] models aspects of the real office in a virtual world, but is based on a virtual reality approach. ClearBoard [19] focuses on the integration of communication and cooperation over a shared drawing tool. Here the integration is deep, resulting in a seamless integration of the inter-personal (communication) space and the shared work space. DIVA in contrast focuses on a broad integration of multiple CSCW aspects. The cooperative Sepia system [14, 15] provides some support for both communication (audio) and cooperation, but is limited to a single application. Several media spaces target synchronous communication, and some include facilities to support informal communication [9, 13, 28].

Polyscope provides background awareness in a work group by periodically distributing low resolution frame-grabbed snapshots of participants' offices to other people's workstations [6]. Portholes extends Polyscope to distributed work groups, using the Internet to transfer the images [11]. These systems convey more information about activities occurring in the physical office space (e.g. co-worker is on the phone, in a meeting, or alone) than does DIVA, but provide fewer clues as to activities occurring amongst distributed collaborating groups (e.g. who is working with whom and on what).

Vrooms [6] is similar to DIVA in that it uses the metaphor of the virtual room and what it means for two or more people to be in the same room—they can see and hear one another. The difference between DIVA and Vrooms is that in Vrooms when two users occupy the same virtual room they share a slow-scan video only connection—as in Polyscope—while in DIVA the connection is by default full-motion video with audio. In Vrooms, a full bandwidth connection is made by moving the two room occupants close together, while in DIVA this action results in the establishment of a private conversation.

### **CONCLUSION**

We have presented DIVA, a virtual office environment for supporting group work. DIVA may be viewed as a prototype replacement for the single user GUI desktop as it is similar in scope but extended to meet the needs of distributed groups—it is the backdrop against which day to day group work is conducted. Initial observation of casual DIVA users suggests that the interface is very easy to learn and use.

The main achievements of this work are, first, the integration of six key CSCW activities within a single working environment. DIVA supports communication, cooperation, and awareness in both the synchronous and asynchronous modes. Second, these services are all provided by a simple and intuitive direct manipulation interface which is both an abstraction of a very common space—the real world office—and an evolution of the GUI desktop. Almost all of the information provided by the system and the operations performed by the users fit smoothly within the office metaphor. Finally, we believe that DIVA introduces several new mechanisms or services within the six CSCW activity classes supported. One such service is the generic support provided for asynchronous merge and catch-up in groupware tools. A second is the awareness information provided in the room view: who is working with whom on what and which documents have changed or are in use by others.

### **FUTURE WORK**

In the near future our development efforts will focus on providing better tools for asynchronous work. In particular we are working on much more sophisticated support both for understanding the changes others have made to a shared artifact (catch-up) and for merging multiple versions of a shared artifact. Both facilities will be extensions to our multi-user GINA application framework and will be based on the manipulation of user activity histories.

While providing catch-up through activity replay is good in the sense that it works for almost all groupware applications, rather than just for text editing, and it provides a consistent interface across applications, it also has problems. The simple replay of changes as they occurred may not be optimal for the viewer. For example the ordering may be poor (often a start-to-finish ordering will be desired), and the granularity may be too fine (character changes, rather than word or sentence level changes). We are working on solving these problems by developing an advanced catch-up facility which provides more than one view of the changes, including views which reorder and coalesce change events, and provides sophisticated animation of the changes, in order to minimize the effort required to understand them.

### **ACKNOWLEDGMENTS**

The system and the concepts described in this paper have been developed by the entire GINA/DIVA project team: Michael Spenke, Thomas Berlage, Christian Beilken, Andreas Genau, and the authors.

## REFERENCES

1. Baecker, R. New Paradigms for computing in the nineties. In *Proceedings of Graphics Interface'91* (June 3-7, Calgary, AB, Canada), CIPS, Toronto, 1991, pp. 224-229.
2. Baecker, R.M., Nastos, D., Posner, I.R., and Mawby, K.L. The User-centred Iterative Design of Collaborative Writing Software. In *Proceedings of CHI'93 Human Factors in Computing Systems* (Apr. 24-29, Amsterdam, The Netherlands), ACM/SIGCHI, NY, 1993, pp. 399-405.
3. Beaudouin-Lafon, M. and Karsenty, A. Transparency and Awareness in a Real-time Groupware System. In *UIST'92, Proceedings of the ACM Symposium on User Interface Software and Technology* (Nov. 15-18, Monterey, CA), ACM/SIGCHI, NY, 1993, pp. 171-180.
4. Berlage, T. and Genau, A. A Framework for Shared Applications with Replicated Architecture. In *UIST'93, Proceedings of the ACM Symposium on User Interface Software and Technology* (Nov. 3-5, Atlanta, Georgia), ACM/SIGCHI, NY, 1993, pp. 249-258.
5. Berlage, T. and Spence, M. The GINA Interaction Recorder. In *Proceedings of IFIP TC2/WG2.7 Working Conference on Engineering for Human-Computer Interaction* (Aug. 10-14, Ellivuori, Finland), 1992.
6. Borning, A. and Travers, M. Two Approaches to Casual Interaction over Computer and Video Networks. In *Proceedings of CHI'91 Human Factors in Computing Systems* (Apr. 28-May 2, New Orleans), ACM/SIGCHI, NY, 1991, pp. 13-19.
7. Buxton, W.A.S. Telepresence: Integrating Shared Task and Person Spaces. In *Proceedings of Graphics Interface'92* (May 11-15, Vancouver, BC, Canada), CIPS, Toronto, 1992.
8. Buxton, W.A.S. Ubiquitous Media and the Active Office. To appear in *Communications of the ACM*.
9. Cool, C., Fish, R.S., Kraut, R.E., and Lawery, C.M. Iterative Design of Video Communication Systems. In *Proceedings of CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 25-32.
10. Dourish, P. and Bellotti, V. Awareness and Coordination in Shared Workspaces. In *Proceedings of the CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 25-38.
11. Dourish, P. and Bly, S. Portholes: Supporting awareness in a distributed work group. In *Proceedings of CHI'92, Human Factors in Computing Systems* (May 3-7, Monterey, CA), ACM/SIGCHI, NY, 1992, pp. 541-547.
12. Fahlen, L.E., Brown, C.G., Stahl, O. and Carlsson, C. A Space Based Model for User Interaction in Shared Synthetic Environments. In *Proceedings of the CHI'93 Human Factors in Computing Systems* (Apr. 24-29, Amsterdam, The Netherlands), ACM/SIGCHI, NY, 1993, pp. 43-48.
13. Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. and Buxton, W. Realizing a video environment: Europarc's RAVE system. In *Proceedings of CHI'92 Human Factors in Computing Systems* (May 3-7, Monterey, CA), ACM/SIGCHI, NY, 1992, pp. 27-35.
14. Haake, A. and Haake, J.M. Take CoVer: Exploiting Version Support in Cooperative Systems. In *Proceedings of CHI'93 Human Factors in Computing Systems* (Apr. 24-29, Amsterdam, The Netherlands), ACM/SIGCHI, NY, 1993, pp. 406-413.
15. Haake, J.M. and Wilson, B. Supporting Collaborative Writing of Hyperdocuments in SEPIA. In *Proceedings of CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 138-146.
16. Hardock, G., Kurtenbach, G. and Buxton, W. A Marking Based Interface for Collaborative Writing. In *UIST'93, Proceedings of the ACM Symposium on User Interface Software and Technology* (Nov. 3-5, Atlanta, Georgia), ACM/SIGCHI, NY, 1993, pp. 259-266.
17. Heath, C. and Luff, P. Disembodied Conduct: Communication Through Video in a Multi-media Office Environment. In *Proceedings of CHI'91 Human Factors in Computing Systems* (April, New Orleans, Louisiana), ACM/SIGCHI, NY, 1993, pp. 99-103.
18. Henderson, D.A., Jr. and Card, S.K. Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, 5, 3 (July 1986), 211-243.
19. Ishii, H., Kobayashi, M. and Grudin, J. Integration of inter-personal space and shared work space: Clearboard design and experiments. *ACM Transactions on Office Information Systems*, 11, 4, (Oct. 1993), 349-375.
20. Johansen, R. *Groupware: Computer Support for Business Teams*. The Free Press, New York, 1988.
21. Kraut, R. and Egido, C. Patterns of contact and communication in scientific research collaboration. In *Proceedings of CSCW'88* (Sept., Portland, Oregon), ACM/SIGCHI, NY, 1988, pp. 1-12.
22. Kreifelts, T., Hinrichs, E., Klein, K.-H., Seuffert, P. and Woetzel, G. Experiences with the DOMINO office procedure system. In *Proceedings of European CSCW'91* (Sept., Amsterdam, The Netherlands), Kluwer, Dordrecht, 1991, pp. 117-130.
23. Leland, M., Fish, R.S., and Kraut, R.E. Collaborative Document Production Using Quilt. In *Proceedings of CSCW'88* (Sept., Portland, Oregon), ACM/SIGCHI, NY, 1988, pp. 206-215.
24. Levergood, T.M., Payne, A.C., Gettys, J., Treese, G.W. and Stewart, L.C. AudioFile: A Network-Transparent System for Distributed Audio Applications. In *Proceedings of the USENIX Summer Conference* (June, 1993).
25. Li, J. and Mantei, M. Working together, virtually. In *Proceedings of Graphics Interface'92* (May 11-15, Vancouver, BC, Canada), CIPS, Toronto, 1992, pp. 115-122.

26. Malm, P.S. The unOfficial Yellow Pages of CSCW. Available by anonymous ftp from the University of Tromso, Norway gorgon.tft.tele.no (192.135.199.112) as pub/groupware/cscwyp.\*, 1993.
27. Mantei, M. Capturing the capture lab concepts: a case study in the design of computer supported meeting environments. In *Proceedings of CSCW'88* (Sept., Portland, Oregon), ACM/SIGCHI, NY, 1988, pp. 257-270.
28. Mantei, M.M., Baecker, R.M., Sellen, A.J., Buxton, W.A.S., Milligan, T., and Wellman, B. Experiences in the Use of a Media Space. In *Proceedings of CHI'91 Human Factors in Computing Systems* (April, New Orleans, Louisiana), ACM/SIGCHI, NY, 1993, pp. 203-208.
29. Medina-Mora, R., Winograd, T., Flores, R., and Flores, F. The Action Workflow Approach to Workflow Management Technology. In *Proceedings of CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 281-288.
30. Neuwirth, C.M., Chandhok, R., Kaufer, D.S., Erion, P., Morris, J., and Miller, D. Flexible Diff-ing in a Collaborative Writing System. In *Proceedings of CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 147-154.
31. Nunamaker, J.F. Dennis, A.R., Valacich, J.S., Vogel, D.R., and George, J.F. Electronic meeting systems to support group work. *Communications of the ACM*, 34, 7, (July 1991), 1991, pp. 30-61.
32. Olson, G.M. and Olson, J.S. User-Centered Design of Collaboration Technology. *Journal of Organizational Computing*, 1 (Jan. 1991), 61-83.
33. Root, R.W. Design of a Multi-Media Vehicle for Social Browsing. In *Proceedings of CSCW'88* (Sept., Portland, Oregon), ACM/SIGCHI, NY, 1988, pp. 25-38.
34. Shneiderman, B. *Designing the User Interface*, Second Edition. Addison-Wesley, Reading, MA, 1992.
35. Sohlenkamp, M. A Virtual Office Environment Supporting Shared Applications. In *Proceedings of the 3rd International Montpellier Conference "Interface to Real and Virtual Worlds"*, (Feb. 7-11, Montpellier, France), EC2, Paris, 1994, pp. 163-173.
36. Spenke, M. and Beilken, C. An Overview of GINA—the Generic Interactive Application. In *User Interface Management and Design, Proceedings of the Workshop on User Interface Management Systems and Environments* (June 4-6, Lisbon, Portugal), Springer Verlag, Berlin, 1990, pp. 273-293.
37. Spenke, M.: From Undo to Multi-User Applications—the Demo. In *Proceedings of CHI'93 Human Factors in Computing Systems* (Apr. 24-29, Amsterdam, The Netherlands), ACM/SIGCHI, NY, 1993, pp. 468-469.
38. Stefik, M. Foster, G., Bobrow, D.G., Kahn, K., Lanning, S. and Suchman, L. Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. *Communications of the ACM*, 30, 1 (Jan. 1987), 1987, pp. 32-47.
39. Wolf, C.G. and Rhyne, J.R. Communication and Information Retrieval with a Pen-based Meeting Support Tool. In *Proceedings of CSCW'92* (Oct. 31-Nov. 4, Toronto, Canada), ACM/SIGCHI, NY, 1992, pp. 322-329.