

REALIZING A VIDEO ENVIRONMENT: EUROPARC'S RAVE SYSTEM

*William Gaver, Thomas Moran, Allan MacLean, Lennart Löfstrand,
Paul Dourish, Kathleen Carter, William Buxton*

Rank Xerox Cambridge EuroPARC
61 Regent Street, Cambridge CB2 1AB, U.K.
gaver@europarc.xerox.com

ABSTRACT

At EuroPARC, we have been exploring ways to allow physically separated colleagues to work together effectively and naturally. In this paper, we briefly discuss several examples of our work in the context of three themes that have emerged: the need to support the full range of shared work; the desire to ensure privacy without giving up unobtrusive awareness; and the possibility of creating systems which blur the boundaries between people, technologies and the everyday world.

KEYWORDS: Group Work, Collaboration, Media Spaces, Multi-Media, Video

INTRODUCTION

Work at EuroPARC involves collaboration among people separated by the architecture of our building and the distance to overseas colleagues at PARC. We have turned this difficulty into an opportunity to research technologies that support collaboration. Many of the most important facets of this work involve the Ravenscroft Audio Video Environment (RAVE). RAVE is an example of a "media space" – a computer-controlled network of audio-video equipment used to support collaboration – which shares features with systems being developed elsewhere (e.g., 9, 19, 23, 25).

In this paper, we focus on three aspects of our research in order to provide an introduction to our media space:

- We want to support shared work over its entire range – from the sort of casual awareness that keeps us informed about the whereabouts and activities of our neighbors to the more focussed and planned work that is involved in joint problem-solving. The current controls of our media space reflect this concern, having evolved with our use of a user-tailorable interface to the system.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1992 ACM 0-89791-513-5/92/0005-0027 1.50

- We are concerned about privacy, but are hesitant about achieving it at the expense of media spaces' ability to provide unobtrusive awareness. We consider the attributes of privacy as many-dimensional. Currently, we combine control and feedback in RAVE to maintain privacy without a loss of functionality.
- We are developing the RAVE system to allow a seamless transition between support for synchronous collaboration and systems which support semi-synchronous awareness over long distances and of planned and electronic events. In this way, we hope to blur the traditional boundaries between people, technologies, and the everyday world, relying both on new technologies and an understanding of people's interactions in the everyday world (cf. 20).

We have been developing a number of systems which use the RAVE infrastructure to enhance our working environment and promote collaboration. In this paper, we discuss examples of systems which have been in relatively wide-spread use at EuroPARC in order to give a taste of the environment we have been developing and to sketch out the philosophy behind this research.




Figure 1: The RAVE system lets us work together in a "media space" as well as the physical workspace.

THE RAVENSCROFT AUDIO VIDEO ENVIRONMENT (RAVE)

EuroPARC was founded in 1987 and there are currently about 30 staff members. Our building, called Ravenscroft House, has 27 rooms and 5 open areas on 4 floors. Despite the small size of the lab, the layout separates us to a surprising degree, so that the building is effectively a collection of relatively isolated sites. One of the motivations for the work described here was to turn this problem into a research opportunity: Because EuroPARC is a small research lab, we were able to install complete data, audio, and video networks throughout the lab. Each room in the building has several audio and video cables running to and from a central switch as well as access to digital networks (see 3 for details). The resulting system, called RAVE, provides all rooms with some form of an audio-video "node," consisting of a camera, monitor, microphone and speakers, which users can move and turn on or off at will. Connections among nodes are completely computer controlled, so that people can display the views from various cameras on their desktop monitors, set up two-way audio-video connections, etc. (see Figure 1). Using this system, we live in a *media space* (25) as well as the physical workspace.

The RAVE system provides us with a great deal of potential functionality. An important design issue concerns how best to constrain it, both to support and encourage its use in ways that enhance existing work practices and to discourage possible misuse (e.g., spying, monitoring, etc.). In considering this question, it is helpful to consider our first design theme, that of supporting the range of collaboration from casual awareness to focussed engagement.

From Awareness to Collaboration

What is collaboration? One perspective – assumed implicitly by much of the current work on CSCW – is of two or more people focussed intensely on a single task. We prefer a broader approach, one we feel better reflects the range of activities involved in shared work. Figure 2 provides a simple representation of our view of what it means to work together.

Two dimensions characterize this framework. The first, *degree of engagement*, refers to the extent to which a shared focus is involved. The second, *amount of planning*, refers to the extent that shared activities occur spontaneously or are planned in advance. Although the space of shared work is probably characterized by many more than two dimensions, this framework allows us to consider four relevant landmarks of the space.

Underlying all is *general awareness*. This simply refers to the pervasive experience of knowing who is around, what sorts of things they are doing, whether they are relatively busy or can be engaged, and so on. Neither planned nor involving a great degree of interaction, this sort of awareness acts as a foundation for closer collaboration – one of the reasons that physical proximity is a highly

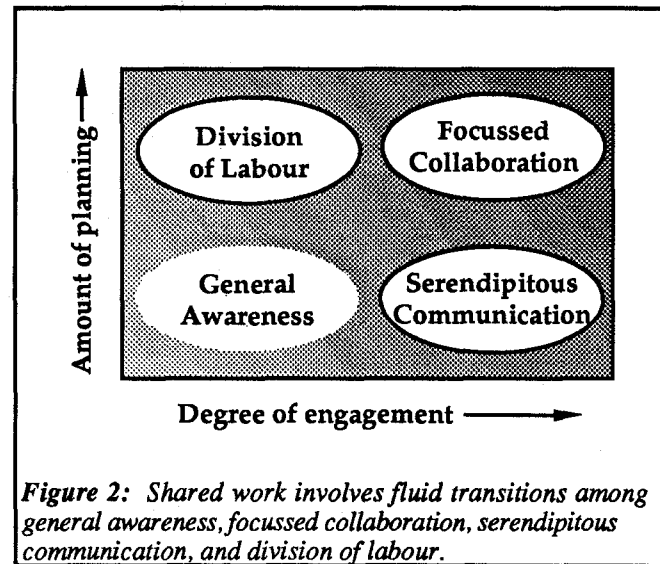


Figure 2: Shared work involves fluid transitions among general awareness, focussed collaboration, serendipitous communication, and division of labour.

accurate predictor of collaboration (15). At the other extreme is *focussed collaboration*. This refers to occasions when people plan to work closely on a shared task. Most CSCW applications seem designed to support this kind of shared, focussed activity.

There are two way-stations between these extremes. The first, *division of labour*, refers to the common practice of splitting a task into its component parts and allowing different people to address them separately. Division of labour does not require the intensely shared focus of attention implied by focussed collaboration, but does require planning and coordination. On the other hand, general awareness often leads to *serendipitous communication*, in which an unplanned interaction may lead to the exchange of important information or the recognition of shared interests.

The description of collaboration illustrated by this framework suggests the need to provide support for a range of activities, from spontaneous to highly planned and from disengaged to highly focussed. Moreover, we want to support the *movement* between these forms of shared work. In the workaday world, people move fluidly between degrees of engagement: maintaining awareness of their colleagues, engaging in serendipitous communication, collaborating intensely for a time, and dividing labour. It is important that we support not only different sorts of shared activities, but fluid movement among activities.

The RAVE Buttons

In providing access to the audio-video network, then, we have emphasized its use in supporting the entire range of shared activities. Because we had few *a priori* notions of how audio-video connectivity would extend current work practices, we have supported access to its functionality in a flexible way, using tailorable onscreen buttons such as those shown in Figure 3.

Buttons are the product of research both at Xerox PARC (14) and at EuroPARC (17). They are simple graphical

objects which allow users to run small programs without having to enter the relevant commands explicitly. In addition, they are tailorable in a number of ways: Their onscreen location and appearance can be modified; they may be copied and emailed; they are often parameterized so that application-specific variables can be changed easily, and their encapsulated code can be edited. Their flexibility allowed us to explore our media space, developing more useful control structures as we gained experience.

Initially, the RAVE buttons provided access to relatively low level functionality, allowing single connections to be made or broken. Over time, the buttons have been modified by users to reflect the higher-level tasks they wished to accomplish. The result is the series of generic RAVE buttons shown in Figure 3.

These buttons reflect the range of engagement in collaboration discussed above – indeed, the buttons and our account of collaboration evolved together. The *background* button, for instance, allows people to select a view from one of the public areas to display on their monitor. This is typically the default connection. Many of us, for example, maintain a view of our largest public space on screen when not actively using the audio-video system. This allows us to notice people come and go to check their mail or get coffee, to see meetings form, or to watch for somebody with whom we want to talk. The effect is similar to having the common area outside one's door (without the noise). We can maintain a general awareness, not of our immediate surroundings, but of important areas that are more remote.

The *sweep* button provides another way to maintain awareness of remote locations of the building. This button makes short (~1 second) one-way connections to various nodes in the building. It is customizable, so one can sweep all nodes or a subset of relevant ones. Typically this is used to find out who is around and what they are doing (cf. 23). The *glance* button, which makes a single 3-second one-way connection to a selected node, allows more focussed attention to particular colleagues. Glances are often used to find out if a particular person is in and whether or not he or she is busy. Because both the sweep and glance buttons allow one way connections for only a

short time, the effect is similar to walking by somebody's door and glancing in: general information about somebody's presence and activities can be obtained without jeopardizing privacy (an issue to which we return below).

More focussed interactions are supported by the *vphone* and *office share* buttons. The first is a two-way audio and video connection which allows colleagues to engage in the video equivalent of a face-to-face conversation. When a vphone call is initiated, the recipient must explicitly accept the connection. Thus this sort of connection is closest to traditional telephone calls. Office share connections are identical to vphone connections, but are meant to last longer – for hours, days, or even months. The effect is one of sharing an office, but because audio volume can be controlled and the video image is relatively small, the other person's "presence" allows but does not demand social engagement.

It is interesting to note here that the vphone and office share buttons offer exactly the same functionality, that of setting up a two-way audio and video connection. The buttons are differentiated solely in terms of the intentions with which the connections are made. Vphone calls are typically used to support relatively short and focussed conversations, while office share connections typically support longer lasting shared work in which the degree of engagement varies fluidly. This is a good example of interface tools which emerged to control our system in terms of users' tasks, rather than technological functionality.

In sum, the five generic RAVE buttons emerged through a process of interconnected use and design supported by an interface system that affords flexible tailoring. The resulting functionality supported by these buttons reflects the range of shared work from general awareness to focussed collaboration to a remarkable degree. The system is even more useful in conjunction with other tools, as we will describe below. But first, it is worth addressing a common set of concerns about the RAVE system.

WHAT ABOUT BIG BROTHER?

Accounts of cameras in every office, one-way glance connections, long-term monitoring of public spaces and so forth can often have Orwellian overtones. Clearly there is a need to protect privacy in audio-video systems such as ours. But there is a trade-off between protection of privacy and provision of functionality that makes the development of such safeguards a non-trivial task.

For instance, one way to assure that our work on media spaces will not add new threats to privacy would be simply to remove all audio and video equipment from EuroPARC – but this would clearly do away with any and all services these technologies offer. More subtly, privacy might be ensured by enforcing symmetrical connections, so that seeing or hearing somebody implies being seen or heard oneself (indeed, this strategy has been taken at BellCore; 23). But one-way connections have advantages we are unwilling to give up. Glances allow us to maintain our




Figure 3: RAVE buttons reflect different degrees of engagement.

awareness of colleagues without actually engaging in interaction with them. Thus they are a valuable prelude to communication; just as we might look in someone's door to see if they are busy before entering, so we can look at their video image before vphoning them. Video provides an excellent means to gain general awareness unobtrusively; enforcing symmetry for the sake of privacy would undermine this functionality.

It has become clear to us that privacy is a complex issue that must be disentangled in order to understand the tradeoffs involved in its protection. In particular, four important facets of privacy which may be considered separately are:

- The desire for *control* over who can see or hear us at a given time;
- The desire for *knowledge* of when somebody is in fact seeing or hearing us;
- The desire to know the *intention* behind the connection; but
- The desire to avoid connections being *intrusions* on our work.

The trade-off between privacy and functionality involves a conflict between the desirability of control and knowledge and the intrusion implied by activities needed to maintain them [cf. 9]. Having to allow explicitly every connection made to our cameras would give us control, but the requests themselves would be intrusive. Having somebody's face appear on our monitors every time they connect to us would similarly demand some sort of social response and might well disrupt previous connections. Having to specify and be informed of the intention of various connections would likewise transform a simple process into a relatively effortful and attention-demanding one. The challenge of safeguarding privacy, then, is not just one of providing control and notification, but doing so in a lightweight and unobtrusive way.

At EuroPARC, our privacy protection depends to a great degree on social convention – indeed, our culture initially provided our only protection. It is assumed that people will use the system with “good” intentions; that is, that they will not seek information with the intent of using it to harm anybody. Simply speaking, we trust one another. At the same time, social convention encourages people to control their own equipment: They are free to turn their camera to face a wall or out a window; they may keep their microphones switched off, and so forth.

We took this initial strategy for several reasons. First, being “willingly naive” about privacy meant that we did not assume the degree to which software support for privacy would be necessary, but instead could treat the question as a research issue. Second, explicitly relying on trust established clear social norms about the use of the

media space – instead of building software on the assumption that privacy would otherwise be invaded, we assumed it would not be and expected people to behave accordingly. Finally, this strategy allowed us to concentrate on developing the functionality of the system rather than security measures. Nonetheless, as the equipment has become ubiquitous in our own lab and we begin to export it to other settings, we have started to explore other ways to tackle privacy issues. Our current system now provides services which make intentionality an implicit feature of connections and which allow us to provide both control and notifications.

Offering Control: Godard

A certain amount of control over connections is offered by the basic software used to control the audio-video switch. This software, called *iiif* (for integrated, interactive, intermedia facility; see 3), instantiates a simple patchbay metaphor in which device “plugs” are linked to form single point-to-point connections. Each plug and device is “owned” by an associated user and its access is accordingly controlled. Thus people could restrict access to their video-out plugs, for instance, to some subset of users.

In practice this strategy is awkward to use effectively. Control is offered at the level of individual connections rather than relevant tasks, while the generic RAVE services described above – glances, vphone calls, etc. – usually involve a number of individual connections. Although buttons can make this transparent to the initiator by combining a number of connection requests into one button, the system has no way of knowing the intention of individual connections. Thus it is difficult, using simple plug control, to design the system so that a glance can be allowed but a vphone call denied.

For these reasons, a new layer of software called *Godard* (7) has been added to the basic *iiif* software. *Godard* uses *iiif*'s underlying protection mechanism to control device plugs so that no connections can be made without its permission. Because *Godard* mediates all connection requests, explicit services can be defined and control can be handled at the service level. When an initiator requests a service, *Godard* uses information previously obtained from potential recipients to determine whether to perform the service (and occasionally relies on interactive input to request permission for individual connections or to resolve conflicts). If permission is given and all relevant plugs are available, *Godard* creates a record of pre-existing connections so they can be restored, and then makes and protects the appropriate connections.

This architecture allows privacy control to exist at the level of services rather than individual plugs. Thus people can set permission for specific people to use specific services. For instance, Figure 4 shows a “glance control panel.” The panel presents a complete list of nodes at EuroPARC, and allows the user to select those who will or will not be given permission to glance. Similar control panels exist for vphones, office share connections, and the like.

With the addition of Godard, our system now affords a degree of *control* adequate to preserve privacy: We can now explicitly allow or deny connections to our equipment. In addition, because these connections are represented as higher-level services, the system also provides a useful (if implicit) representation of the initiator's *intentions*. Finally, it serves as a foundation for the provision of the third aspect of privacy suggested above, that of *knowledge* of actual events – notifications about the system state.

Providing Notifications: Auditory Cues

Feelings of privacy are not only supported by control over who can connect to one's equipment using various services, but by feedback about when such connections are actually made. Because Godard knows about connections to recipients' audio-video nodes at the service level, it facilitates the provision of such feedback. Several kinds of feedback can be requested by users in current instantiations of interface software, including text messages displayed on their workstations and spoken messages played over the audio network. Less obvious than these, and in our experience quite valuable, are *auditory cues* used to provide information about system state (11).

For example, when a glance connection is made to a camera, Godard triggers a sound to be played at the relevant location (the default is that of a door opening). The sound typically comes several seconds before the connection is actually made, so it provides forewarning rather than concurrent information. When the connection is broken, another sound (typically that of a door closing) is triggered. In addition, different sounds indicate different sorts of

connections (and thus the intentions behind them). A knock or telephone bell indicates a vphone request; door sounds indicate glances; footsteps might indicate sweeps; and a camera whir indicates that a framegrabber has accessed one's node. Thus auditory cues provide information about what kind of connection is being made, over and above information about the existence of a connection alone.

Playing sounds such as opening and closing doors may seem frivolous, but nonspeech audio as a medium has several advantages over graphics, text or speech:

- Sound indicates the connection state without requiring symmetry – that is, it provides information without being intrusive.
- Sounds such as these can be heard without requiring the kind of spatial attention that a written notification would.
- Non-speech audio cues often seem less distracting and more efficient than speech or music (although speech can provide different sorts of information, e.g., *who* is connecting).
- Sounds can be acoustically shaped to reduce annoyance (22). Most of the sounds we use, for instance, involve a very gradual increase in loudness to avoid startling listeners.
- Finally, caricatures of naturally-occurring sounds are a very intuitive way to present information. The sound of an opening and closing door reflects and reinforces the metaphor of a glance, and is thus easily learned and remembered (cf. 12).

These sorts of auditory cues have provided an flexible and effective way to unobtrusively inform people that somebody is connecting to their node, and thus serve as another means of safeguarding privacy. More generally, with Godard and auditory cues, we have provided control, feedback, and intentionality – three prerequisites for privacy – at very little cost in terms of intrusiveness. Big Brother would have a difficult time at EuroPARC, both because we can restrict his access and because we can hear him coming.

AWARENESS OVER TIME: THE KHRONIKA SYSTEM

Our audio-video system has helped us maintain awareness of ongoing events in distant locations. *Khronika* (16) is a software "event notification service" that supports selective awareness of *planned* and *electronic* events. *Khronika* is related to online calendar systems, but supports a more general notion of events than most. It tells us when a video connection has been made, reminds us about upcoming meetings, provides information about visitors, and can even be used to gather people to go to the pub.

Khronika is based on three fundamental entities: *events*, *daemons*, and *notifications* (see Figure 5). Events are defined in terms of their class, their start time, and their




Figure 4: Control panels allow users to give permission to specific individuals for specific services.

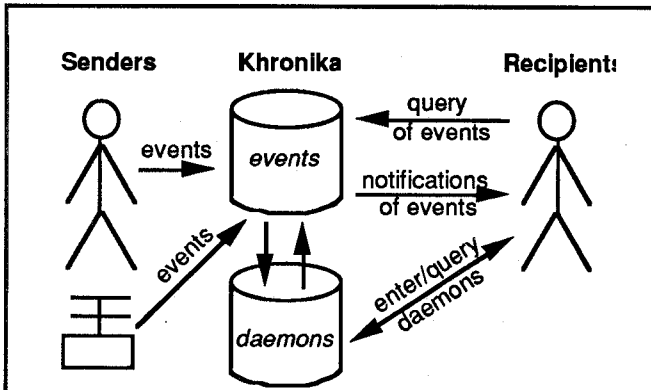


Figure 5: Khronika maintains a database of events entered both by people and other systems. Daemons watch for specified events and post notifications when they are detected.

duration. Examples of events include conferences, visitors, local movies, and arriving email. Because they are represented as objects in a hierarchical classification structure, they can also be manipulated in terms of more abstract classes such as "professional," "electronic," and "entertainment."

Event daemons watch for specified event types and produce notification events when they are detected. Daemons are created by users as a set of constraints, so recipients choose the information about which they wish to be informed. For example, a user may create a daemon which watches for all seminar events occurring in the conference room with the string "RAVE" as a part of their description. They can then instruct the daemon to generate notifications

five minutes before relevant seminars are due to begin.

A number of interfaces to the Khronika system have been explored, including buttons which allow users to browse the event database and to create new events and daemons. One of the more interesting and useful interfaces is the *xkbbrowser*, shown in Figure 6. The browser serves as an online calendar, with events shown as fields extending over their relevant times. But the event database may be displayed at varying levels of specificity, from the most encompassing ("event") level to more specific ones such as "meetings," "glances" or "sound." In this way, the *xkbbrowser* provides a general and powerful mechanism for exploring the database of events.

Notifying Users About Events

Khronika is the mechanism with which Godard generates feedback about audio-video connections. When a request for a connection is made, Godard enters an event into Khronika; an appropriate daemon (created using the various privacy controls already described) then triggers the requested notification.

Notifications can be generated by daemons in several different forms – for instance, a daemon watching for meetings might send out an email message the day before, display a message on a workstation window, or generate a synthesized speech message. Nonspeech audio cues are commonly used to inform us about the state of the audio-video system; there are also a number of cues which inform us about other events (see 11).

For example, we are often reminded about upcoming meetings by the sound of murmuring people gathering together, followed by a gavel sound. This sound acts as a

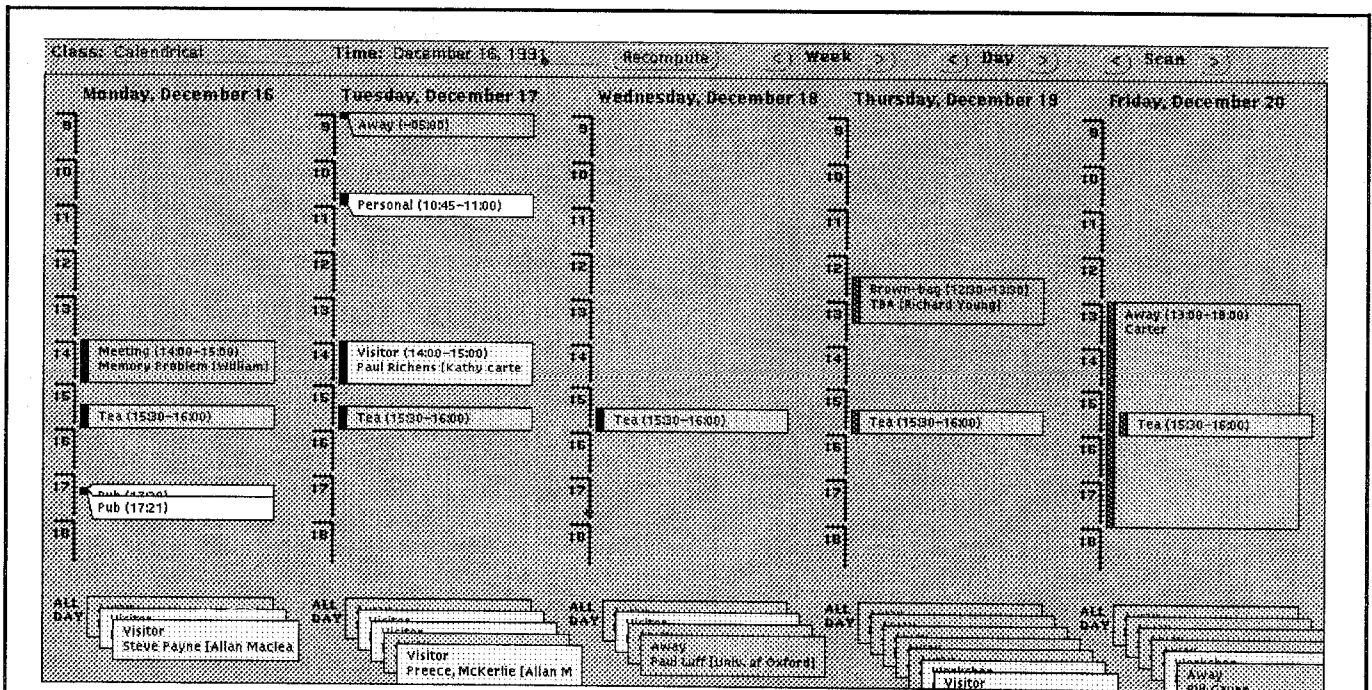


Figure 6: The *xkbbrowser* lists events in a calendar-like format. Event types can be seen at various levels of specificity.

memorable stereotype of naturally-occurring meeting sounds and is thus quickly learned and immediately recognizable. In addition, the sound is designed so that it grows in amplitude quite slowly, so that is not interruptive. Finally, the sharper gavel sound at the end lends a sense of urgency to the sound. Sounds like these are effective yet unobtrusive reminders about remote events – as evidenced by the fact that approximately 50 sounds a day are requested from the Khronika system.

In general, then, the Khronika system in conjunction with audio reminders has a number of the system features we are exploring at EuroPARC. It enhances our general awareness of ongoing events and thus promotes collaboration. It does so in a way that blurs the boundaries between the electronic and everyday worlds, allowing information to be entered from and disseminated by both. Finally, it allows for a great degree of user customization and, like all our systems, is in a continual state of evolution guided by use.

AWARENESS OVER SPACE: POLYSCOPE AND PORTHOLEES

RAVE is useful in providing awareness of local nodes. But for technical and financial reasons, we cannot make connections to our overseas colleagues, nor can we connect to more than one node at a time. In order to extend our awareness over a greater distance and to a number of people simultaneously, we have been experimenting with distributing low-resolution video images via our digital networks.

An initial prototype, *Polyscope* (2) is a system which we used to distribute digitized images within our building every 5 minutes or so. The resolution of the images is not very high – only 200 by 150 bits, with no grey scale. Nonetheless, people and objects in their environments are usually visible. In addition, a simple animation facility is available, in which a few images are digitized successively and looped on display. Although such animations are often jerky (and sometimes deliberately frivolous, as when one researcher arranged to periodically transmogrify into Elvis Presley), they make movement obvious and are an effective way to disambiguate scenes. Moreover, *Polyscope* acts as an interface to the audio-video network. Buttoning an image produces a pop-up menu which allows glance or vphone connections to be initiated.

We are currently using a more recent version of this kind of system called *Portholes* (8). The major advantage of *Portholes* over *Polyscope* is that it runs between EuroPARC and PARC – this means that we can see images of colleagues in a building about 6,000 miles away with those of people in our own building. Not only does this support awareness, but it has helped to create and develop a new research community within EuroPARC and PARC – for instance, researchers who have never been co-present nonetheless speak of "knowing" one another through their experience with *Portholes*.

Both *Polyscope* and *Portholes* allow several remote locations to be presented simultaneously, affording passive awareness of distributed workgroups without the necessity of explicitly setting up video links and so on. This facilitates smooth transitions between general awareness and more focussed engagements. In addition, the spatially-distributed but asynchronous functionality offered by systems like *Portholes* and *Polyscope* complements our synchronous but single-channeled video services quite well. Perhaps most importantly, *Portholes* allows us to extend this awareness out of our building to colleagues at geographically distant locations.

EXPERIENCE, EXPERIMENTS, AND EXPORT

We have said little about our experiences using these systems. In general, our development efforts rely on what might be considered a form of participative design, in which designers work closely with users in shaping useful systems (4). At EuroPARC, as with most research labs, the division between designers and users is often blurred. Nonetheless, the group can be divided into technical and non-technical staff, and much of our development is guided by the experiences and input of non-technical users (see 17 for an example of this process). In addition, a number of users have been keeping diaries of their experiences with various systems. These accounts are a valuable source of insight about audio-video mediated collaboration.

More formal techniques have also been useful in better understanding the nature of our media space. Ethnomethodological and participative design techniques have been employed to study the everyday use of the RAVE system and to assist in its development. For example, observations of video-mediated communication have indicated that the medium can undermine the effectiveness of subtle communicative gestures (13), leading us to explore ways to enhance our system. In addition, a series of open-ended interviews have been used to identify problems with the system as well as new possibilities for its design (5).

We have also used more traditional experimental studies to examine a range of issues. For instance, a recent study assessed the utility of a collaborative text editor called *ShrEdit* and the effects of shared video on its use (21). Another study examined patterns of gaze associated with task and meta-level conversations among co-located or remote partners working in a shared software environment (24). In a third study, we found that nonspeech audio feedback changed participants' perception of a complex collaborative system and their tendency to collaborate while using it (10).

Finally, we have begun exporting these technologies to new sites to better understand how they interact with and support existing work practices. For example, recent research on participative design has involved the installation of a limited audio-video link in a London architecture firm (6). Building on this, a new project is using audio-video technologies to support designers

working together but based in different countries – England and the Netherlands (18).

REALIZING A VIDEO ENVIRONMENT

In this account we have been concerned with describing RAVE and several of the related systems we use to support shared work at EuroPARC. We have suggested ways these systems work together to form an integrated environment, and have sketched some of their philosophical foundations.

We hope to have given a feeling for the kinds of systems we are developing. Moreover, we hope to have shown that the three themes of our research – supporting the range of collaboration, maintaining privacy, and extending media spaces to include awareness of planned, electronic, and semi-synchronous events – provide a valuable foundation for research on collaborative systems which are integrated across the working environment. Above all, we have tried to convey a sense of why we find the research at EuroPARC fun, exciting, and important.

ACKNOWLEDGEMENTS

Our work at EuroPARC has depended on the collaborative efforts of many people. In particular, Bob Anderson, Ian Daniel, Christian Heath, Paul Luff, Tom Milligan, Wendy Mackay, Mike Molloy, Toby Morrill, Gary Olson, Judy Olson and Randall Smith have all contributed to the development of the research described here, as have our colleagues at PARC Sara Bly, Steve Harrison, Austin Henderson, Scott Minneman, and John Tang. More generally, the entire EuroPARC research community has provided invaluable support for this work simply by using these systems as part of their everyday environment. Finally, the surrounding philosophy has been emerging from our research community at PARC and EuroPARC for years; it is impossible to assign credit for most of these ideas.

Tom Moran's current address: Xerox PARC, 3333 Coyote Hill Drive, Palo Alto, CA. 94304.

Bill Buxton's current address: CSRI, University of Toronto, Toronto Ontario, Canada M52 1A4

REFERENCES

- Bellotti, V., Dourish, P., & MacLean, A. (1991). From users themes to designers DREAMS: Developing a design space for shared interactive technologies. EuroPARC/ AMODEUS Working Paper RP6-WP7.
- Borning, A., and Travers, M. (1991). Two approaches to casual interaction over computer and video networks. *Proceedings of CHI'91* (New Orleans, Louisiana, 28 April - 2 May, 1991). ACM, New York, pp. 13-19.
- Buxton, W., and Moran, T. (1990). EuroPARC's integrated interactive intermedia facility (iiif): Early experiences. In *Proceedings of the IFIP WG8.4 Conference on Multi-User Interfaces and Applications* (Herakleion, Crete, September 1990).
- Carter, K. (forthcoming). Interacting with users: A practitioner's experience. To appear in *Sociology of Software*, Woolgar, S., and Murray, F. (eds).
- Carter, K. (July, 1991). Usage of the AV network. Presentation at EuroPARC RAVE Review, Cambridge U.K.
- Carter, K. and Harper, R. (1991). Searching for problems and answers: An empirical report on CSCW. Technical Report No. EPC-91-101, Rank Xerox EuroPARC, Cambridge U.K.
- Dourish, P. (1991). Godard: A flexible architecture for AV services in a media space. Technical Report No. EPC-91-134, Rank Xerox EuroPARC, Cambridge U.K.
- Dourish, P., and Bly, S. (1991). Portholes: Supporting awareness in a distributed work group. *Proceedings of CHI'92* (Monterey, California, 3 - 7 May, 1992). ACM, New York.
- Fish, R., Kraut, R., Root, R., and Rice, R. (1991). Evaluating video as a technology for informal communication. *Proceedings of CHI'92* (Monterey, California, 3 - 7 May, 1992). ACM, New York.
- Gaver, W. W., Smith, R. B., and O'Shea, T. (1991). Effective sounds in complex systems: The ARKola simulation. *Proceedings of CHI'91* (New Orleans, Louisiana, 28 April - 2 May, 1991), ACM, New York.
- Gaver, W. W. (1991). Sound support for collaboration. In *Proceedings of ECSCW'91* (Amsterdam, The Netherlands, 25-27 September 1991).
- Gaver, W. W. (1986). Auditory icons: Using sound in computer interfaces. *Human-Computer Interaction*, 2, pp. 167-177.
- Heath, C., and Luff, P. (1991). Disembodied conduct: Communication through video in a multi-media office environment. *Proceedings of CHI'91* (New Orleans, Louisiana, 28 April - 2 May, 1991), ACM, New York.
- Henderson, D. A., and Card, S. (1986). Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, 5, 3, 211-243.
- Kraut, R. and Egido, C. (1988). Patterns of contact and communication in scientific research collaboration. In *Proceedings of the CSCW'88* (Portland, Oregon, September 1988) ACM, New York. pp. 25-38.

16. Lövstrand, L. (1991). Being selectively aware with the Khronika system. In *Proceedings of ECSCW'91* (Amsterdam, The Netherlands, 25-27 September 1991).
17. MacLean, A., Carter, K., Moran, T., and Lövstrand, L. (1990). User-tailorable systems: Pressing the issues with Buttons. In proceedings of CHI'90 (Seattle, Washington, 1-5 April, 1990) ACM, New York, pp. 175-182.
18. Mackay, W., and Harper, R. (1991). WAVE: The Welwyn and Venray Experiment. Technical Report No. EPC-91-135, Rank Xerox EuroPARC, Cambridge U.K.
19. Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T., and Wellman, B. (1991). Experiences in the use of a media space. *Proceedings of CHI'91* (New Orleans, Louisiana, April 28 - May 2, 1991) ACM, New York, pp. 203 - 208.
20. Moran, T. P. and Anderson, R. J. (1990). The workaday world as a paradigm for CSCW design. In *Proceedings of CSCW'90* (Los Angeles, California, October 1990). ACM, New York.
21. Olson, G., and Olson, J. (1991). User-centered design of collaboration technology. *Journal of Organizational Computing, 1*, 61-83.
22. Patterson, R. D. (1989). Guidelines for the design of auditory warning sounds. *Proceedings of the Institute of Acoustics 1989 Spring Conference. 11, 5*, 17-24.
23. Root, R. W. (1988). Design of a multi-media vehicle for social browsing. In *Proceedings of the CSCW'88* (Portland, Oregon, September 1988) ACM, New York, pp. 25-38.
24. Smith, R. B., O'Shea, T., O'Malley, C., Scanlon, E., and Taylor, J. (1989). Preliminary experiments with a distributed, multi-media, problem-solving environment. *Proceedings of the First European Conference on Computer-Supported Cooperative Work* (Gatwick, England).
25. Stults, R. (1986). Media space. Xerox PARC technical report.