

ONE IS NOT ENOUGH: MULTIPLE VIEWS IN A MEDIA SPACE

William Gaver^{1,4}, Abigail Sellen^{1,2}, Christian Heath^{1,3}, Paul Luff^{1,3}

¹ Rank Xerox Cambridge EuroPARC
61 Regent Street, Cambridge CB2 1AB, UK
<surname>@europarc.xerox.com

³ Department of Sociology
University of Surrey, Guildford GU2 5XH, UK

² MRC Applied Psychology Unit
15 Chaucer Road, Cambridge CB2 2EF, UK

⁴ Industrieel Ontwerpen
Technische Universiteit Delft
Jaffalaan 9, 2628 BX Delft, Nederland

ABSTRACT

Media spaces support collaboration, but the limited access they provide to remote colleagues' activities can undermine their utility. To address this limitation, we built an experimental system in which four switchable cameras were deployed in each of two remote offices, and observed participants using the system to collaborate on two tasks. The new views allowed increased access to task-related artifacts; indeed, users preferred these views to more typical "face-to-face" ones. However, problems of establishing a joint frame of reference were exacerbated by the additional complexity, leading us to speculate about more effective ways to expand access to remote sites.

KEYWORDS

cscw, social interaction, media spaces, video

INTRODUCTION

Over the last few years, several laboratories have built experimental *media spaces*, computer-controlled networks of audio and video equipment, in order to explore issues concerning the support of collaboration among distributed colleagues [4, 6, 12, 18]. Media spaces provide more than basic video conferencing and picture-phones, creating a "space" for interaction that exists alongside the everyday physical environment. They are intended to support a wide range of shared work activities, from focused collaboration to a more general awareness of events in the workplace [6].

Despite the promise of these systems, our experience suggests that their main value is in providing casual awareness of remote colleagues or as a prelude to more focused, co-present communication [see also 4]. Current systems *can* support focused collaboration, but it seems that even experienced users prefer to be physically co-present for a variety of focused collaborative tasks.

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.

© 1993 ACM 0-89791-575-5/93/0004/0335...\$1.50

Experimental studies, on the other hand, tend to suggest that allowing people to see one another does not add significantly to the process of collaboration. Comparisons among co-present, video plus audio, and audio only interactions suggest that access to visual information has no significant effect on the dynamics of conversation [14, 15]. Visual access also tends not to affect the outcome of intellectual, decision-making, and creative tasks. Tasks that rely on social cues, in contrast, such as situations of conflict, bargaining and negotiation do tend to be affected by the presence or absence of the visual communication channel [1, 13, 16].

Does this mean that media spaces might support general awareness, but are no better than telephones for focused collaboration? In this paper, we suggest that the problem is not inherent to the medium, but rather to the beliefs about collaboration that guide the design of current media spaces. In particular, we challenge the assumption that a face-to-face, head and shoulders view is most important for collaboration – an assumption made both in most media spaces and in most experiments assessing the effects of video [though see, e.g., 19]. We describe a study in which we evaluated the possibility of overcoming the limitations of media spaces by providing multiple views on a collaborative task. Using these data, we discuss the strengths and weaknesses of the multiple camera strategy and speculate about more effective ways to offer increased access to remote sites.

FACE TO FACE VIEWS ARE NOT ENOUGH

Face-to-face interaction is only one of many shared activities undertaken by co-present collaborators, yet most media spaces are configured to support little else. For instance, Figure 1 shows a typical media space node: a video monitor sits on the desk, with a camera on top of it and a microphone and speakers somewhere nearby. This sort of arrangement allows local users to see their colleagues, but prevents access to many other aspects of the remote environment, such as their documents, their workstations, and other nearby objects. Although cameras may be moved, in practice they seldom are, because of the inconvenience of repositioning cables, the jerky images such movement produces, and the difficulty of aiming the camera at the desired location.



Figure 1: A typical media space node

Observational studies of naturally-occurring collaboration suggest that many of the subtle interactions involved in successful shared work are disrupted by the limitations of these kinds of views. Workers in control rooms and architectural practices, for instance, perform much of their work without explicit face-to-face communication [11]. Instead they modify their speech and gestures to provide and acknowledge information, revealing considerable peripheral sensitivity to their colleagues' activities. The limited views provided by typical media spaces disrupt these ways of working [9]. Collaborators in media spaces are often frustrated by their inability to show each other artifacts such as paper or screen-based documents [see also 4]. Moreover, many of the actions performed via nonverbal behaviour seem to lose their interactional significance when conducted through a media space. In sum, current media spaces fail to support peripheral monitoring and involvement in collaborative tasks.

These observations are complemented by an analysis of the affordances of media spaces – the properties of the environment that have implications for perception and interaction [5, 7]. Perhaps the most important limitation of the audio-video medium, according to this analysis, is that cameras are stationary or only moved locally in most media spaces. This restricts the possibility of visually exploring remote spaces, and implies that information must be actively presented to remote colleagues. It also seems to worsen problems associated with the asymmetry of information flow typical of media spaces. For instance, seeing a colleague in a media space does not imply that one can be seen. Prediction is difficult, and people cannot move to compensate for their partners' inaccuracies.

Taken together, experimental, sociological, and ecological studies converge in suggesting significant constraints due to the limited field of view offered by a single fixed camera configured for face-to-face interaction. These constraints go beyond the inability to share access to a relevant artifact. Rather, collaboration depends on the ability to explore perceptually the relevant task environment, maintaining a sensitivity towards one's colleagues in the context of their current activities. This does not imply that face-to-face

views are useless, but it does stress the need to provide more, and more flexible, access to remote sites.

THE MULTIPLE TARGET VIDEO STUDY

As a first attempt to provide more flexible access to remote working environments, we designed an experimental system called Multiple Target Video (MTV) which allowed users to switch among multiple cameras in a remote site. We then observed pairs of collaborators using the system to work on two shared tasks.

The MTV System

Two offices at EuroPARC were equipped with four cameras each as well as two monitors, a microphone, and speakers, and connected with audio and video links (see Figure 2). In each office, the view of the remote office was shown on a colour monitor and selected using a video switch with a rotary control. In addition, whatever the person in the remote office was viewing was shown on a smaller, monochrome "feedback monitor".

We arranged the cameras to provide views based on our understanding of the visual information that co-present individuals use during shared work. One camera in each office was mounted over the main monitor, as it is in typical media spaces, to provide a *face-to-face view*. A second camera was mounted above and to the side of the main work area to give an *in-context view* of people in relation to their workspace and the objects within it (including the two video monitors). The third view was provided using a high-resolution, monochrome camera, giving a *desk view* allowing documentation to be shared.

These views were roughly the same for each office; however, the fourth camera was arranged differently depending on the office. In the *Design Office*, the camera was arranged to support one of the tasks that we gave participants. This task, which will be described below, involved the manipulation of a physically embodied architectural representation (i.e., a dollhouse) located in the office. Thus the fourth camera in the Design Office provided a *dollhouse view* for the remote participant. In the *Remote Office*, a camera was mounted in a corner of the room and a wide-angle lens was used to provide a *birds-eye view* showing most of the room.




Figure 2. The MTV set-up.

Data Collection

Further connections were made to a third office between the two, allowing the study to be monitored and data to be collected. We used a picture-in-a-picture device to record four video streams simultaneously. The in-context view from each office was recorded to provide an overall impression of the participants' activities and the views that participants were currently using were recorded to observe their switches and the information they had available.

The Study

Six pairs of participants took part in the study: two from EuroPARC, and four from the MRC Applied Psychology Unit subject pool. Members of most pairs had known each other for at least a few months, but one of the pairs had not met before the study. The two pairs recruited from EuroPARC had some experience with media spaces, while the other four had not.

The Room Task

The first task was designed to familiarise the participants with the equipment and with each other. Each participant was asked to draw the layout of their partner's office, including major pieces of furniture and equipment. We explicitly suggested that they help one another, because neither of the offices was entirely visible from the combination of views they had available. Subjects were told that it should take about 10 or 15 minutes to complete the task, and that they would be observed and recorded from the third office.

The Design Task

In the second task, we asked participants to select and arrange furniture in a dollhouse located in the Design Office, to draw the final layout, and to write a list of reasons for their decisions. This task was designed to encourage collaborative work that focused on and around relevant artifacts. One member of each pair was assigned to the Design Office containing the dollhouse and thus could move the dollhouse furniture directly. The other member of the pair was in the Remote Office, and had to rely on their partner to rearrange the furniture.

Participants were given conflicting hidden agendas for the task. One was assigned the role of a designer from the Cheep 'n' Tasteless Vacation Rental Co., and told to imagine that he or she would receive a substantial bribe if all the furniture were used. The second participant, in contrast, was given the role of a designer from the award-

winning Integrity Design Company and told to arrive at a tasteful solution that did not necessarily use all the furniture. In addition, Cheep 'n' Tasteless designers were told to finish the design quickly because they were paid for the job, while Integrity designers were told not to hurry, both to make sure that the design was optimal and because they were paid by the hour. Participants were not explicitly told that they would be rivals, but most seemed aware that some sort of conflict had been arranged. They usually took about 45 minutes to complete the task, but we stopped them if they seemed unable to reach a solution.

The aim of these tasks was two-fold. First, previous studies have largely neglected video-mediated collaboration centered around objects [though see 8]. We wanted to evaluate the MTV system's ability to support such collaboration, particularly when the objects were physically available to only one partner. Second, previous empirical results [1, 13, 16] have suggested that face-to-face interaction is important for negotiating conflict. We expected the focus on objects to reduce face-to-face interaction, but wanted to see if people would use face-to-face connections more in the competitive Design Task than in the cooperative Room Task.

RESULTS AND DISCUSSION

All the participants were able to work together successfully on the tasks, and the multiple cameras clearly supported them in doing so. Both tasks were designed to be highly visual in nature, so it is difficult to imagine that participants could have pursued them using face-to-face or audio connections alone. MTV was more complex than typical face-to-face media spaces, but we found that participants quickly developed a range of strategies for using the views to pursue the tasks.

Viewing patterns are captured by strip charts such as the one shown in Figure 3, which shows the switching behaviour of a representative pair of participants over an entire session. The charts show the five possible views that participants might have (face-to-face, in context, desk, birds-eye, and doll house) vertically, and record the time each participant spent in a view along the horizontal axis. This representation provides a general impression of participants' switching behaviour. For instance, this pair tended to switch more in the Room Task than they did in the Design Task, a pattern we found for all participants. In addition, relatively long periods using one view were punctuated by short bursts of switching among the others.

Pair 4




Figure 3. Switching data from an entire session. Hollow strips show the view from the Design Office; solid strips that from the Remote Office. Only the Design Office had the birds-eye view; the Remote office had the dollhouse view.

Table 1: Percentage of Time in Each View

task	office	face/face	con-text	desk view	birds eye	doll-house
room	design	6	12	32	50	-
	remote	14	58	27	-	1
design	design	12	24	36	28	-
	remote	10	12	7	-	71
TOTAL		11	23	24	17	25

As can be seen from Figure 3, all the cameras were used by the participants. Table 1 shows the percentage of times spent in each view averaged over all participants. Random viewing would yield percentages of about 25%. Instead, clear patterns of use become evident from these data.

Face-to-Face Views Were Rarely Used

One of the most striking results is that participants rarely used the face-to-face views (only 11% of the total time). Even more rarely were they truly "face-to-face," with both members simultaneously using the face-to-face views (about 2%). Instead, the switching times and video data suggest that participants preferred to maintain views which showed objects more directly relevant for the tasks.

The only substantial period of shared face-to-face viewing came about when one pair of participants engaged in lengthy negotiations near the end of the Design Task. Most participants maintained a shared focus on relevant artifacts (i.e., the dollhouse or drawing) even while engaged in intense discussions. In fact, several stated that they avoided using the face-to-face views, particularly when they were trying to convince their partner of some action.

These data suggest that, given the choice between face-to-face views and views that give access to shared work objects, people will choose the latter. This is not due merely to the conflict we introduced in the Design Task; the face-to-face views were seldom used in either task. It also did not appear to be affected by the familiarity of partners; people who had never met before behaved in much the same way as those who knew each other well. Instead, it seems to be because both tasks were designed to require access to the remote environment and objects within it, as suggested by observations of real-world collaboration [11].

The assumption inherent in the design of many current media spaces that face-to-face views are the most valuable is seriously challenged by these data. This is not to suggest that face-to-face views are useless. We saw numerous occasions in which participants would make very short connections to their partners' face-to-face view, seemingly to assess their engagement and mood. The numerical data here may well under-emphasize the importance of such "glances," but it also highlights the utility of other views.

Use of Views

The video data suggest a variety of ways that participants used the views. For instance, the face-to-face view was used as much to see the office behind the remote partner as to engage in conversation. The in-context and birds-eye views were used to gain access to the remote office for the Room Task, to determine the remote partner's current orientation, and to observe pointing at video images or at the room itself. The desk view was used to offer information to partners about one's local office, to check sketches of the remote one, and to provide detailed views of artifacts such as the dollhouse furniture. Finally, the ability to switch between views provided flexible access to the remote office. For instance, it allowed one participant to track her partner moving around the office pointing out features of interest. In general, each view was not used merely for one predictable purpose; rather, the views were appropriated as resources for pursuing the tasks.

Limitations of MTV

While there were clear benefits of the increased access to remote environments offered by multiple cameras, difficulties with the approach became evident upon analysis of the video data. Understanding the problems of using multiple cameras to extend access to remote sites is valuable in suggesting alternative possibilities for design.

Lost in Space

In the following sections, we describe a representative episode in some detail to illustrate the difficulties participants had while working with the MTV system. This episode occurred as a pair of participants were working on the Room Task. Mary, attempting to draw Sue's office, asks about a detail she cannot see. Sue starts to sketch her own office to show Mary, but has problems making clear which part of the room she is sketching:

Sue: ... No, none of the monitors are going to show you any of it. I tell you what, I'll sketch it for you...

Mary: Oh brilliant. Yeah, great.

Sue: Now behind me, right...

(Mary switches to desk view to see Sue's sketch.)

Sue: If that's the corner of the room... (pointing to part of her sketch) where you're getting the view of me...

(Sue reaches towards her video switch, pauses.)

Sue: Let me just see what... Just switch to the view of me... (Mary switches to birds-eye view.) Right. (Sue points to feedback monitor, which shows Mary's view of her.) Can you see that, the corner where... there's a sort of... (stops pointing at monitor, points up to corner of the room) ..I'm pointing up to the corner now.

Mary: Yeah.

Sue: Right... uh going back along that wall... (reaches for her video switch, withdraws) Now go back to my drawing.

(Mary switches through several views to the desk view)

Mary: Hang on... there you are.

Sue: There we are, right... (drawing on page) There's that... and then the doorway is like this...

Consider how easy it would have been to establish the correspondence between the drawing and the room if the two partners had been co-present. A mere gesture over the document followed by one towards the relevant part of the room would have served. The participants in this episode, in contrast, used a variety of methods to achieve the same goal, facing continual errors and breakdowns of reference. While they managed to recover from these problems, the greater access provided by the MTV system was clearly balanced by its limitations.

Each of the problems that these participants encountered is representative of those facing all the users of this system. In the following sections, we describe each of these difficulties in turn.

The Configuration is Never Right

Sue: No, none of the monitors are going to show you any of it. I tell you what, I'll sketch it for you...

Despite the use of four cameras per office, each adjusted to provide what we felt was a maximally effective view, it became clear on observing participants that the views provided were still not optimal for the tasks we gave them.

Evidence for the inadequacy of views can be found not only in participants' remarks, but in their use of the views as shown in Figure 3 and Table 1. For each task, participants spent a clear majority of time using only one of the four views. In the Room Task, participants tended to use the views that gave them the most access to the remote site – the in-context view for those in the Design Office, and the birds-eye view for those in the Remote Office. In the Design Task, inhabitants of the Remote Office were completely dependent on the views for access to the dollhouse. Not surprisingly, most used the dollhouse view. This view allowed them to monitor the current arrangement of furniture, but it did not allow them to keep track of their remote partner. In fact, one remote participant used the in-context view most during the Design Task, claiming afterwards that she took this strategy so that she could keep track of both her partner and the dollhouse. Inhabitants of the Design Office, in contrast, usually seemed to ignore the video connections, concentrating instead on the dollhouse itself. Some even made use of the limited access their partner had: one participant argued that her partner couldn't judge the arrangement because of the poor view provided by video, while another hid a piece of furniture simply by moving it off camera.

Observations such as these suggest that no configuration of stationary cameras will provide adequate access for all tasks. In the Room Task, participants often compensated for the limited access they had by drawing their own offices for their partners, as Sue does in the example above. This raised further difficulties, however, in establishing the relations between the drawing and the room itself.

Difficulties in Establishing Reference

Sue: Now behind me, right...

(Mary switches to desk view to see Sue's sketch.)

Sue: If that's the corner of the room... (pointing to sketch) ...where you're getting the view of me...

The different views provided by the cameras caused problems in relating different scenes to one another – for instance, in pointing out a feature of the sketched layout and then indicating which part of the room it represented. There were problems in mapping one view to another, and in establishing the correspondence between a representation (e.g., a drawing) and its referent.

Difficulties in establishing reference have been noted before in descriptions of traditional point-to-point media space connections [9]. By increasing the number of cameras, however, we seem to have seriously increased these problems. Not only did participants have to cope with the differences between their view of a remote colleague and their colleague's view of them, but with discontinuities in the views they had of the remote office, differences in their view of a local artifact and that of their partner, and so forth. Strategies for dealing with these problems were further constrained by difficulties in establishing mutual orientation, as discussed below.

Unequal Control Of Orientation

(Sue reaches towards her video switch, pauses.)

Sue: Let me just see what... Just switch to the view of me... (Mary switches to birds-eye view.) Right.

The system also hindered the ability to direct the attention of remote partners, interfering with the establishment of a mutual orientation to relevant objects. In co-present situations, people use gestures, pointing, and so forth to direct their colleagues' attention and negotiate a mutual orientation towards a task. The ability to do this might have helped in establishing the relations among various views, but these tactics were ineffective using the MTV system. Indeed, one-sided control over the views meant that even less subtle tactics failed: participants seemed to want control over their partners' views, with aborted gestures towards their own switches often preceding a verbal request that their partners change views.

Problems involving control over orientation and action arose in other situations as well. For instance, co-present partners can gesture to indicate where marks should go on a document if they don't have a marker. The MTV system made this difficult. We observed several participants try to direct their partners by pointing at a document on their monitor while their partner watched using the in-context view, but this appeared to have limited success. More often participants would resort to verbal directions. For instance, one pair of subjects working on the Room Task had difficulty in establishing shared reference to a drawing:

John: Now this is hard... how do I point at your document? Ann: I know, I also had trouble – I was pointing to the TV screen... which is pretty stupid, because you can't see it.... How about if I move my finger

around and then you can tell me when to stop? (Ann moves finger over drawing.) John: Ok, up and to the left... (Ann moves finger accordingly)...

In both of these situations, participants' inability to direct attention flexibly to things both in the local and remote environments meant that they had to resort to explicit verbal control. Though a number of strategies were developed to deal with these problems, all were relatively troublesome.

Video Images as Shared Artifacts

Sue: (points to feedback monitor, which shows Mary's view of her) Can you see that, the corner where... there's a sort of... (stops pointing at monitor, points up to corner of the room) ..I'm pointing up to the corner now.

Because people had unequal access to objects themselves, video images sometimes seemed to serve as shared artifacts to a greater degree even than the local scenes they displayed. In the Room Task, for instance, we observed many occasions in which local participants, trying to describe their own offices to their participants, would point to features on the feedback monitor (which showed the view seen by their remote partner) rather than to the room itself. Similarly, inhabitants of the Design Office frequently checked their feedback monitor to assess the current appearance of the dollhouse while arguing for some arrangement. In these instances, participants appeared to act towards images as objects having greater mutual accessibility than the things they represented.

Treating video images as shared artifacts caused problems, however, when participants assumed that they had access not only to the image, but also to the space around it. Co-present collaborators communicate using gestures and movements within the space around a shared artifact [2], but this was impossible using the MTV system. We repeatedly witnessed participants gesturing and pointing at the screen image of the document as if the space around that document was mutually accessible. These gestures, of course, largely went unnoticed by their partners; even when they were seen using the in-context or birds-eye views their gestures did not move in the space around the artifact, but had to be related to the artifact by their partners. This inequality of access to artifacts compounded the problems of control and reference discussed above.

Switching is Problematic

*Sue: ...Now go back to my drawing.
(Mary switches through several views to the desk view.)*

Mary: Hang on... there you are.

Sue: There we are, right...

The necessity of switching among views increased the difficulty of establishing the relations among them and negotiating a mutual orientation towards the task. The arbitrary mapping between the rotary switch and the views it provided was difficult to learn, making the results of a

given switch effectively unpredictable. Participants seemed to avoid the need to remember this mapping – or indeed the existence of various views – by quickly switching through the views looking for one that was adequate. Though this was usually satisfactory, views sometimes seemed to be forgotten when a scan did not reveal them. Additional experience with the system might alleviate this problem, but only if a stable set of views and settings were used.

More seriously, the restriction to seeing only one view at a time prevented a smooth transition of awareness to objects in different parts of the remote office – participants could not shift easily between the drawing and the remote office, for instance. The limited access participants had to their partners' activities disrupted the ability to negotiate mutual orientation fluidly; instead changes of orientation were often signalled only by a change on the feedback monitor. In general, the constraint implied by switched views contributed to many of the problems that participants faced. For this reason, we plan to replicate this study in the future, replacing the video switch with separate, continuously available monitors for each view.

CONCLUSIONS

In sum, providing multiple cameras allowed participants to collaborate successfully on two tasks involving remote artifacts. Moreover, our data clearly indicated that participants preferred task-centred over face-to-face views. While face-to-face views are undoubtedly useful for some tasks, these results suggest that other views provide a wide range of resources for remote collaborators.

However, our observations of the MTV system in use suggest a number of difficulties in using multiple cameras to provide access to remote sites:

- Configuring cameras to give complete access to a complex space is difficult or impossible.
- Multiple views interfere with establishing a shared frame of reference.
- One-sided control over views and artifacts inhibits negotiations of mutual orientation.
- Unequal access to the space around shared artifacts poses problems of coordination.
- Switching among multiple views accentuates their discontinuities.

Understanding these difficulties is helpful in suggesting new possibilities for extending media spaces. For instance, access to video images could be supported by providing telepointers to users, similar to those found in many shared drawing systems [see 3] but using video overlay technology [e.g., 10]. Participants could be given equal control over views by allowing them to switch views on their own office as well as remote ones. Finally, the incorporation of shared drawing tools might overcome the problems in establishing a shared frame of reference for, and mutual orientation to artifacts. However, we suspect similar difficulties would remain, given that the need to

share physically embodied artifacts is likely to remain a feature of shared work.

Many of the difficulties that participants experienced using the MTV system seem to come from the need to switch between multiple, discontinuous views of remote spaces. As an analysis of the affordances of media spaces [5] suggests, the ability to move continuously in remote spaces might help with these problems. Movement allows us to change our focus of attention smoothly from place to place, either on our own initiative or in response to a colleague. This enables us to interactively establish a mutual frame of reference towards objects of interest.

From this perspective, many of the problems in using the MTV system may be ascribed to the fact that it did not allow true movement in remote sites. Techniques that allow continuous camera movement in remote spaces might overcome the limitations we observed in this study. Tracking cameras, for instance, which automatically follow a subject as he or she moves around an area, would allow participants to follow their partners' actions, but at the expense of self-guided exploration. Pan-tilt-zoom units, already used in some video-conferencing systems, would allow self-guided movement – but such systems are often slow and awkward to use. Perhaps most promising are systems that allow intuitive control over remote movement, such as the Delft depth television system [17]. The incorporation of technologies such as these into future versions of media spaces is a promising possibility.

Despite its drawbacks, however, the MTV system clearly represents a first step towards extending the capabilities of media space systems. This is an important goal: If media spaces are to serve as a true alternative space for interaction, we must question basic assumptions about how they are designed and used.

ACKNOWLEDGEMENTS

We thank Wendy Mackay for useful discussions about the design of the system and study reported here. Anne Schlottmann suggested using the dollhouse, and she, Steve Harrison, and Allan MacLean made helpful comments on this paper. Finally, Mike Molloy helped design the MTV system and set it up single-handed. For any errors in this paper, each author cheerfully blames the others.

REFERENCES

1. Argyle, M. and Cook, M. *Gaze and mutual gaze*. London: Cambridge University Press, 1976.
2. Bly, S., and Minneman, S. Commune: A shared drawing surface. *Proceedings of COIS'90*, ACM, New York, 1990, pp. 184 - 192.
3. Dourish, P., and Bellotti, V. Awareness and coordination in shared work spaces. *Proceedings of CSCW'92* (Toronto, 1 - 4 November, 1992).
4. Fish, R., Kraut, R., Root, R., and Rice, R. Evaluating video as a technology for informal communication. *Proceedings of CHI'92* (Monterey, 3 - 7 May, 1992). ACM, New York, pp. 37 - 48.
5. Gaver, W. The affordances of media spaces for collaboration. *Proceedings of CSCW'92* (Toronto, 1 - 4 November, 1992).
6. Gaver, W., Moran, T., MacLean, A., Lövsstrand, L., Dourish, P., Carter, K., and Buxton, W. Realizing a video environment: EuroPARC's RAVE system. *Proceedings of CHI'92* (Monterey, California, 3 - 7 May, 1992). ACM, New York, pp. 27 - 35.
7. Gibson, J. J. *The ecological approach to visual perception*. Houghton Mifflin, New York, 1979.
8. Harrison, S., Minneman, S., and Irwin, S. Graspable implication: A study of 3-D objects in remote collaboration. Xerox PARC technical report, 1992.
9. Heath, C., and Luff, P. Media space and communicative asymmetries: Preliminary observations of video mediated interaction. *Human-Computer Interaction*, 7 (1992), pp. 99 - 103.
10. Ishii, H., and Kobayashi, M. ClearBoard: A seamless medium for shared drawing and conversation with eye contact. *Proceedings of CHI'92* (Monterey, California, 3 - 7 May, 1992). ACM, New York, pp. 525 - 540.
11. Luff, P., Heath, C., and Greatbatch, D. Tasks-in-interaction: Paper and screen based documentation in collaborative activity. *Proceedings of CSCW'92* (Toronto, 1 - 4 November, 1992).
12. Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T., and Wellman, B. Experiences in the use of a media space. *Proceedings of CHI'91* (New Orleans, April 28 - May 2, 1991) ACM, New York, pp. 203 - 208.
13. McGrath, J. *Groups: Interaction and performance*, Englewood Cliffs, NJ: Prentice-Hall, 1984.
14. Sellen, A.. Speech patterns in video-mediated conversations. *Proceedings of CHI'92* (Monterey, California, 3 - 7 May, 1992). ACM, New York, pp. 49 - 59.
15. Sellen, A. Statistical measures of speech in technology-mediated conversations. Technical Report No. EPC-92-125, Rank Xerox EuroPARC, 1992.
16. Short, J., Williams, E., and Christie, B. *The social psychology of telecommunications*. London: Wiley & Sons, 1976.
17. Smets, G. Designing for telepresence: The interdependence of movement and visual perception implemented. *Proceedings of the IFAC Man-Machine Symposium*, 1992.
18. Stults, R. Media space. Xerox PARC technical report, 1986.
19. Weber, K., and Minneman, S. The Office Design Project. *SIGCHI Video Review*, 1988.