Culture and Control in a Media Space

Paul Dourish
Rank Xerox EuroPARC, Cambridge, UK
dourish@europarc.xerox.com

Abstract: Media spaces integrate audio, video and computer networking technology in order to provide a rich communicative environment for collaboration. The connectivity which they provide brings with it important concerns regarding privacy, protection and control. In order to derive the fullest benefit from this technology, it is essential that these issues be addressed. As part of our investigation of media space systems, we developed a computational infrastructure addressing these problems our own working environment. A key aspect of this work is the relationship between two aspects of this control system—the *technological* components which determine how the system will behave, and the *social* components which determine acceptable use and behaviour.

This paper discusses our experiences with the privacy and control aspects of our RAVE media space environment, specifically with regard to connection management, and compares them to the experiences of other research groups. We discuss the nature of the relationship between technological and social elements in using this technology, and discuss the consequences for the design of such systems.

Introduction: Media Spaces

One focus for research into workgroup communication and collaboration is the investigation of media space technology. A media space is formed by the combination of audio, video and computer networking technologies to provide a flexible, dynamic interconnection environment for collaborative work groups. Our "RAVE" media space (Buxton and Moran (1990); Gaver *et al* (1991)) is one of a number of systems which have been used in investigations into issues of workgroup support and collaboration (*e.g.* Root (1988); Stults (1989); Abel *et al* (1990); Mantei *et al* (1991)).

At its most basic, the media space provides a means for setting up multi-media communication channels between individuals and groups; beyond that, it supports collaborations, both formal and informal, providing a communication infrastructure which is amenable to rapid reconfiguration and connection. This element of *flexible* and *dynamic* control by individuals is at the heart of what we mean by media space; a media space is more than simply point-to-point audio and video connections. While fixed video links can provide communication between remote individuals, the easy reconfiguration and control of connections is critical in exploring the means by which the technology can create a space in itself, rather than being simply a restructuring of the physical space.

Clearly, the provision of this sort of media infrastructure is a useful means towards contact and direct collaboration between individuals who are not co-located in time or space. The media space is not just a desktop video-conferencing system; our research (*e.g.* Dourish and Bly (1992)) has pointed to other, less formal mechanisms which can be built in media space environments to support a variety of work groups.

On the other hand, it is equally clear that the introduction and use of media space technology raises a host of issues concerning the protection of individual privacy and access control. This is particularly true in environments such as RAVE which are used not only by research staff involved in the development of the technology, but also by other office workers throughout our lab. Even within our own environment, then, it was critical that these issues be addressed if our work was to be successful. The result is a software component called "Godard" (Dourish, 1991) which provides inhabitants of the media space with flexible control over the degree of access they grant to others, and dynamic information on connections and connectivity.

However, control and accessibility in the RAVE environment is not purely based on a technological infrastructure, but also on the culture which has developed around this system. We would claim that any technological control system like Godard sits within a culture which determines aspects of its use; and that this relationship between technological and social control must be considered as part of the design process. In essence, the relation between the technology and its use is *coadaptive*, and must be analysed as such (Mackay, 1990). This relationship is the central focus of this paper.

In the rest of this paper, we will first discuss aspects of the Godard system in more detail, and describe the way in which its introduction changed aspects of media space use within RAVE. We will then consider the design of some other media space environments and aspects of privacy control and usage culture in those systems. Finally, we will talk more generally about the relationship between technological and social control and its importance in designing such systems.

The RAVE Media Space Environment

The Ravenscroft Audio-Video Environment is a media space being developed and used at Rank Xerox EuroPARC. It comprises switchable analogue video and audio networking technology, workstation software to control connections among individuals and public spaces, and other systems which provide facilities for coordination, informal interaction and focussed collaboration. RAVE extends these with computer-based systems which augment the power of the media space for collaboration. Multi-user drawing surfaces (e.g. ShrEdit (McGuffin and Olson, 1992) and shdr (Dourish, *to appear)*), a shared active calendar (Lövstrand, 1991) and other tools for collaboration in a number of computational domains are all important parts of the system as a whole.

One critical aspect of the RAVE system is its ubiquity within our work place. Rather than being restricted to a small group of individuals who are involved in the development of the technology, RAVE connections are provided to *all* members of the lab, including research, technical and administrative staff. Although there is no obligation to make use of the technology, the hardware and associated software systems are made available by default. Furthermore, users are encouraged to "adopt" the technology and make it their own. For instance, there is no obligation to have cameras pointed directly at work areas. Pointing the camera out of a window for privacy during a meeting or conversation is quite common.

Clearly, the ubiquity of our system raises a number of issues concerning privacy and protection for individuals. Being deployed through our lab, the RAVE system does not require the "sign-up" process observed in some other environments, where an individual must explicitly request membership of the media space, and accept whichever norms govern its use. Hence it was extremely important for us to explicitly address issues of privacy and control in the development of our system. It was critical to the success of our experimentation that the participants could maintain the level of privacy they desired, maintain control over the ways in which others could connect to their offices, and have flexible control over feedback indicating the state of the media space and their visibility within it. At the same time, it was also important that we maximise the flexibility and utility of our media space, in order to investigate the range of ways it could support collaborative work in our environment. RAVE is, after all, a research tool.

These requirements led us to develop a solution based on people's styles of interaction in the real world (Anderson and Moran, 1990). This is realised primarily in a system component called Godard.

Godard

Godard is a software infrastructure which provides the primary interface for users of the RAVE media space. It encompasses subsystems which are responsible for making connections and controlling the media space, as well as for controlling access from other individuals and connection status feedback.

There are two concepts which underlie the design of Godard. The first is the notion of a *service*. A service is the basic “unit of interaction” in the RAVE media space. The network offers a range of services to users of the system, and these services are encapsulations of stereotypical user behaviours in a physical office environment. Some examples will give the flavour of what we mean by a service:

1. The glance service provides the media space equivalent of glancing into a colleague’s office as you walk by their door. In the media space, this takes the form of a one-way video connection of a few seconds’ duration.
2. The vphone service is the media space equivalent of a telephone call. A two-way audio and video connection is set up between two individuals after the “recipient” of the connection has answered a ring signal.
3. The office-share service is the media space equivalent of sharing an office with a colleague. A two-way audio and video connection links individuals and can be maintained for as long as both parties agree (some office-share connections have been in place almost continuously for over a year). The individuals can choose to use open (“hands-off”) audio, or to control it with a foot-pedal.
4. The background service is the media-space equivalent of the view out an office door or window. It is a video connection from an office to a public area such as a corridor, the coffee room or even a view out of a window towards the park behind our building. This view will generally be shown on a monitor when it is not engaged in any other activity. It is automatically disconnected when other connections are made, and reconnected after they are over.

The difference between the vphone and office-share connections is particularly interesting, and illustrates the orientation of these services around *user behaviours* rather than technical descriptions of system activity. Both vphone and office-share link offices with two-way audio-and-video connections; they are technologically almost identical. However, their mode of use is very different. Whereas a video phone call will typically be short and directed (another property it has in common with “real” phone calls), an office-share connection between two individuals is normally of much longer duration and much less continuously attended.

This definition of services around stereotypical real-world behaviours, then, gives us a useful way to partition the space of possible connections. By tagging a video connection with some notion of an analogous real-world behaviour, we can provide a way for users to discriminate between connections. For example, users may choose to accept glance connections but not office-share ones. This discrimination is based on an understanding of the kind of connection being made (the service). Technically, it happens through the action of the other principal Godard component, a RAVE *agent*.

The agent is essentially a gatekeeper, which responds semi-autonomously to requests for audio or video connections made by other users. The agent, then, embodies individual privacy within the system. RAVE participants cannot be glanced at unless to appropriate agent has been consulted. The agent decides whether or not to accept a particular connection by consulting a set of individual preferences supplied by the user. These record the details of which connections should be accepted, which should be rejected, and which should cause the agent to interrupt the user for confirmation.

As well as providing a basic level of protection, we can also use this mechanism to provide *feedback* to users on the state of connections. Thus, for instance, a user can control three feedback parameters for the glance service:

1. the Before parameter specifies some action which should take place before the actual video connection is made;
2. the After parameter specifies a behaviour which should take place when the video connection is broken;
3. the Inform parameter specifies a way that the system should inform the user of the *identity* of the individual making the connection.

Typically, this feedback takes the form of non-speech audio cues (Gaver, 1991), or, less often, pop-up messages on the workstation screen. Thus, the before action might generate the sound of an opening door, and the after action may sound like the door closing again. Since these actions are stored individually for each user in a particular agent process, a user has complete control over what form these feedback messages should take, or of they should be used at all, for their own node. Parameter settings for various services remain private, and inaccessible to other users of the system except through their externally-observable result. Since the introduction of our system, the audio cues have proved to be the most popular—probably because they do not require that the user is attending to the screen when a connection is made.

Introducing Godard

A natural question at this point is, in what ways did the introduction of these mechanisms change the nature of the media space, and why? The first answer is that, as had been hoped, privacy protection *made the system more open*. Many more people were prepared to make themselves accessible to their colleagues via the media space for three main reasons:

1. They now trusted that a system was in place to protect their privacy if they wanted, which reduced the cost of changing accessibility. It was possible to make oneself accessible without the worries about how one could undo that again.
2. The presence of feedback in the system meant that users felt they could be part of the system without the original feeling that they didn't know what was going on. Feedback made the status of media space actions much more immediate.

3. The use of different services to delimit the space meant that users could make themselves *selectively accessible*. They could be accessible for short-term connections (glance) without being available for longer-term connections (vphone or office-share), or even long-term one-way connections (watch connections, which are generally only offered for public areas). This gave them a much finer degree of control than had previously been available.

It might be regarded as curious that a system designed to enforce privacy and protection results in a more open communicative environment and increases access. However, this is really quite natural. The openness and accessibility of an environment is not purely a function of the technology, but of the culture of use which arises around it. Godard enabled a different culture of use to emerge. We will consider the point in more detail later, after a brief look at other media space systems and their approaches to these issues.

Other Systems

In the introduction, we referred to a variety of research projects investigating media space technologies. Each of these environments has some means for controlling access, whether by technological or social means. In this section, we discuss some other systems in more detail, with emphasis on the ways in which they control access and availability of participants.

Xerox PARC: Media Space

The system at Xerox's Palo Alto Research Center was probably the earliest example of a media space (in fact, it's called "Media Space"). The use of media space technology at PARC has its roots in a split-site experiment, when part of the System Concepts Laboratory was divided between Palo Alto, California, and Portland, Oregon. The group was linked with a digital audio and video connection, initially between two public areas, but later between two local switching hubs, giving them the opportunity to make inter-office connections between sites (Olson and Bly, 1991). Although the Portland lab has now closed, the Media Space continues as an active research area centered in the Information Sciences and Technologies Laboratory in Palo Alto, and links around 16 offices as well as public areas and shared resources.

The PARC media space operates on a "sign-up" basis. New users wishing to join the media space approach the existing group and must explicitly request access. Part of the process of "signing up" is an acceptance of the social practices and norms which govern acceptable media space use.

There is a minimal technological protection system, which protects inter-node connections already in place from being broken by others, rather than protecting nodes. However, this "locking" mechanism is rarely used. The result is that any

Media Space user can make an audio and/or video connection to any other, and disconnect most ongoing connections. Privacy can be maintained by pointing the camera out of a window, or turning off the microphone.

PARC's protection model, then, is largely a cultural one. While the technology makes it possible for individuals to "misuse" the technology, there are social pressures which prevent them from doing so, and these can be preserved through the "sign-up" model. Within a small community, the result is a stable situation, comfortable and acceptable to participants, without direct need for a more technological solution.

University of Toronto: CAVECAT

The CAVECAT system is a media space under development at the University of Toronto. Experiences of the use of this media space have been reported elsewhere (Mantei *et al*, 1991), but here we concentrate in particular on some of the privacy and accessibility concerns, especially with reference to the other systems under discussion.

The basis of the CAVECAT system is the IIIF connection management system (Milligan, 1991), which is the same system as underpins RAVE. Extensions to manage accessibility and resource contention, which RAVE encapsulates in Godard, are realised at Toronto by extensions to IIIF and intelligent clients, which together form the CAVECAT system (Louie *et al*, to appear).

CAVECAT users can create rules for connections, similar in spirit to the service access lists which Godard maintains. However, these rules are managed differently in practice, since they are based on *explicit properties* of the underlying media connections, rather than on distinguished services. A rule in the CAVECAT system might say:

> **If** one of the users [*tom, marilyn, bill, gifford*] requests a
> connection with properties [*short-term, one-way, video*] to
> my node, **then** accept with notification [*audio, knock*]

This rule refers to glance-type behaviour. Certain common rule patterns are reified as accessibility states presented to users through an iconic interface which draws on the metaphor of office doors. Full accessibility is suggested by an open door; a door which is ajar makes me less accessible, requiring my permission for some sorts of connections; a closed door requires more permission again, and a locked or barred door indicates that the user is completely inaccessible.

The connection properties are superficially similar to Godard's services, but differ in the degree to which they can be distinguished. Godard's services are based on descriptions of real-world behaviour; CAVECAT's are based on technological descriptions of the underlying media connections. This means that Godard can distinguish between connections which are technologically identical but have different modes of use; this is not available in CAVECAT. On the other hand, the CAVECAT rules generalise to new connection types in a way in which Godard's cannot.

One characterisation of this distinction is in terms of the assignment of responsibility which the systems imply. In Godard, the burden of responsibility for declaring the connection type (the service or "context" of connection) lies with the initiator, who has to select the service being requested. In CAVECAT, on the other hand, discrimination takes place on properties of the connection itself; there is no declaration of intent on the basis of the initiating party. Recipients must configure their own rules so as to correctly intuit the mode of use of the system based on properties of the underlying connections. It must be recognised that a long term, two-way video connection from a particular person, for instance, implies a particular mode of use, and hence is to be dealt with in a particular way. The burden of responsibility is on the recipient to form the appropriate generalisations and then structure these into rules.

So, CAVECAT strives to manage protection and control largely through technological measures. In their rules, users create embodiments of the social conventions which mediate their interactions, describing them in terms of connection properties; but however much the rules are a formal technical system, they still record socially-derived preferences for initiating and managing interactions. However, the formal structure of the rules distances users from the social situations they are describing; and the orientation of the rule system puts the burden on responsibility for connection interpretation on the recipient. Managing the social nature of media space interaction in a framework such as this can be difficult, as it attempts to use the technological mechanisms to replace, rather than augment and preserve, the social components which are the essence of media space use.

Bellcore: CRUISER

Another workplace media space is that developed at Bellcore, called "Cruiser" (Root, 1988). Cruiser is designed to provide a mechanism for "social browsing", encouraging informal interactions between work group members, including those separated by distance. The original Cruiser design was based on a hallway metaphor; initiating a "cruise" would cause brief connections to a number of individuals in succession, similar to walking down a hallway and looking into the offices. Either party in a connection could halt the cruise and engage the other in conversation. Later versions of Cruiser added a specific "glance" facility, rather like a one-stop cruise.

In terms of protection mechanisms and privacy control, two aspects of Cruiser are of interest to us here. The first is that all connections in Cruiser are *reciprocal*; that is, when I see another media space participant, she sees me at the same time. It is simply not possible to make a one-way connection with the Cruiser system. The second point is that the members of the Cruiser media space have a simple way of indicating their level of accessibility through the use of computer-generated bars which appear on their video image. The bars obscure the video image of the participant, indicating to the remote observer that this individual is not accessible for conversation.

Cultural aspects of Cruiser use are embodied in its strong “corridor browsing” metaphor, and are carried across from those real-world behaviours. Certain behaviours are “anti-social” (*e.g.* repeatedly “cruising” the same individual), and reciprocity in Cruiser serves something of the same role as feedback in RAVE by making aspects of the system-internal behaviour accessible in the real world and amenable to social control. A price is paid in terms of the potentially increased intrusion of a two-way connection. In fact, a major privacy concern expressed by users focussed not so much on the violation of privacy through connection, but through the imprecise and unpredictable nature of the technology (Fish *et al*, 1992).

A number of technological features of Cruiser, such as reciprocity, essentially serve to support a socially-controlled mechanism for interaction and initiating of conversations. To an extent, then, reciprocality and a strong face-to-face metaphor make it very easy to carry across behaviours from real-world social interaction into Cruiser; but on the other hand, it makes it much more difficult to have new behaviours develop within the media space, since the modes of interaction are strictly limited.

The Social-Technical Continuum

The preceding sections have discussed a number of experimental experiences with the use of media space technologies in work group collaboration. We have shown the way in which different groups have addressed issues of privacy and control, using both social and technological mechanisms to regulate access and accessibility.

One characterisation which we can make of the various privacy controls is their position along a posited “social-technical continuum”. At one end, we might place PARC’s system in which control is largely social; the technology provides many facilities which are not allowed by the social norms which determine its use. RAVE’s Godard solution is close to the other end, in that control (at least on a per-connection basis) is largely technical; far fewer “misuses” are allowed by the technology.

But to what extent is the Godard mechanism purely based in technology? In fact, it would be false to claim that Godard implements a purely technical model of “correct video behaviour” which obviates the need for social controls. For instance, if I grant only glance access to other users, I clearly intend this to provide short-term views into my office; the technology will ensure that requests for long-term connections are denied. However, what is to stop another user from making repeated glance connections to my office and thus watching me over a much longer period of time?

In the RAVE environment, the use of the feedback mechanism deters this type of misuse. Making repeated glance connections in the media space becomes rather like pressing your ear to a colleague's closed door in an office; while the technology does not prevent it, it is easily discovered and clearly against the prevailing culture. So, one of the roles of connection feedback in Godard is to "lift" aspects of control from the technical domain to the social one. The feedback from repeated glance connections means that this form of misuse is amenable to social, rather than technological, control. This relationship between the two modes of control is critical to Godard's success, and a fundamental aspect of its design.

We would claim that even those media space control systems which based themselves strongly in technology are also subject to the culture of use which emerges from them. There are both positive and negative aspects of this cultural embedding. On the positive side, we would simply point out that media space technologies do not necessarily impinge upon and nullify the social conventions which regulate workplace behaviours; the media space is just as amenable to such forces as are physical spaces. The negative side is that the negative aspects of cultural embedding are still present; social pressures may mitigate against, for instance, refusing video access to superiors. (In Godard, one user cannot examine another's agent preferences; however, a superior can still *demand* access as much as before.)

This is not a question of exchanging one evil for another, however; it is merely another example of the way in which social controls still matter in media spaces. In our research environment, Godard is not proposed as *the* way to manage privacy and control issues; rather, it is our particular embodiment in a technical system of the controls which we choose for our workplace. On a more general level, it points towards the importance of the balance between technical and social constraints in these systems.

The raising of control issues into the social domain may also change the way the technology is used. The use of audio feedback in Godard's glance service provides an example of this. Within our environment, it is typical to allow anyone to glance as long as full feedback is enabled. This convention has allowed a different use of glance to arise, in which users may initiate glance connections when they arrive in the morning to say "hello" to their colleagues; they are making use of the cultural convention that a glance connection is announced at the other end.

These cultural phenomena arise every day in normal workplaces. A mundane example is to be found in the conventions that give meaning to open and closed office doors. We are all familiar with the conventions of our own workspaces, and the messages that open and closed office doors send to passers-by about the accessibility of the individual within. One door at EuroPARC, although typically closed, carries a notice which reads "this door is OPEN"; the owner of that office wishes to distinguish between the physical state of the door and the statement of a closed door in our office culture. Similar conventions surround the use of phones, answering machines, voice mail and forwarding mechanisms; letting my phone calls forward to a receptionist has a different meaning from letting them forward to an answering machine.

In other words, the tension between the technological and social barriers which ensure individual privacy and control over access is something which we deal with regularly. It is a tension which is natural to us in our everyday lives. It seems only natural, then, that we should capitalise upon this tension in the design of media space environments. Indeed, it is a contention of this paper that it is impossible not to; technological systems cannot be designed which are not then subject to the influences of social factors in their use.

The social factors which influenced the design of Godard continue to be an important aspect of further design within our environment. In particular, some colleagues are currently investigating the way that low-level information about the connection status of the media space and the use of equipment can be provided in order to reduce the problem of disembodiment of the individual from the communicative environment which characterises aspects of media space interaction (Bellotti and Sellen, 1993). This work pursues the theme of the relationship between social and technical control over privacy and accessibility in media space environments.

Summary

The use of multimedia communication environments, or media spaces, is of increasing interest to research groups studying various forms of collaborative and group working. One critical aspect of these systems is the extent to which individual privacy is maintained and accessibility controlled. If this technology is to be successful in supporting collaboration, it is first essential that means are provided allowing people to control the extent to which they are accessible, preferably without a steep learning curve.

A variety of mechanisms, technological and social in nature, can be used to make the media space a comfortable place to live and work. By discussing mechanisms that various systems use and the ways in which they differ, we suggest that a purely technical notion of protection and control is not only inappropriate, but impossible. Interaction in a media space is by its very nature a social activity, and the technological systems we might use to manage media space connectivity are embedded within social and cultural contexts.

We would argue that an acknowledgment of these social and evolving elements surrounding interaction in media spaces is a critical element in the design of a system to manage privacy and accessibility. This can greatly benefit the technological designer, since it serves, to an extent, to delimit the scope of any technological control; and it allows us to integrate media space interaction much more easily into everyday behaviour by exploiting our everyday understanding of the cultural elements of workplace interaction.

ACKNOWLEDGEMENTS

The development of the RAVE media space has involved contributions from many individuals. I would like to acknowledge Sara Bly, Alan Borning, Bill Buxton, Kathy Carter, Ian Daniel, Mik Lamming, Lennart Lövstrand, Hugh Messenger, Wendy Mackay, Allan MacLean, Tom Milligan, Mike Molloy, Tom Moran, Mike Travers and Alex Zbyslaw for their help and contributions.

The ideas expressed in this paper, and the form of the paper itself, have similarly benefitted from discussions with a number of colleagues, including Bob Anderson, Victoria Bellotti, Sara Bly, Bill Gaver, Steve Harrison, Rachel Jones, Wendy Mackay and Abigail Sellen.

References

Abel, M., Corey, D., Bulick, S., Schmidt, J. and Coffin, S. (1990): "The US West Advanced Technologies TeleCollaboration Research Project", in Wagner (ed.), *"Computer Augmented Teamwork"*, Van Nostrand Reinhold.

Anderson, R. and Moran, T. (1990): "The Workaday World as a Paradigm for CSCW Design", in *Proc. ACM Conference on Computer-Supported Cooperative Work CSCW'90*, Los Angeles, Ca., October 1990.

Bellotti, V. and Sellen, A. (1993): "Designing for Privacy in Ubiquitous Computing Environments", in *Proc. European Conference on Comuter-Supported Cooperative Work ECSCW'93*, Milano, Italy.

Buxton, W. and Moran, T. (1990): *"EuroPARC's Integrated Interactive Intermedia Facility (IIIF): Early Experiences"*, in *Proc. IFIP Conference on Multi-User Interfaces and Applications*, Herakleion, Crete, September 1990.

Dourish, P. (1991): "Godard: A Flexible Architecture for A/V Services in a Media Space", EuroPARC Technical Report EPC-91-134, Rank Xerox EuroPARC, Cambridge, UK.

Dourish, P. and Bly, S. (1992): "Portholes: Supporting Awareness in Distributed Work Groups", in *Proc. ACM Conference on Human Factors in Computer Systems CHI '92*, Monterey, Ca., May 1992.

Dourish, P. (to appear):, "Anatomy of a Shared Workspace: Dissecting Shdr", to appear in Dourish (ed.), *"Implementation Perspectives in CSCW Design"*, Springer-Verlag, in preparation.

Fish, R., Kraut, R., Root, R. and Rice, R. (1992): "Evaluating Video as a Technology for Informal Communication", in *Proc. ACM Conf. Human Factors in Computing Systems CHI '92*, Monterey, Ca., May 1992.

Gaver, W. (1991): "Sound Support for Collaboration", in *Proc. European Conference on Computer-Supported Cooperative Work ECSCW '91*, Amsterdam, September 1991.

Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K.and Buxton, W. (1992): "Realising a Video Environment: EuroPARC's RAVE System", in *Proc. ACM Conference on Human Factors in Computing Systems CHI '92*, Monterey, Ca., May 1992.

Louie, G., Mantei, M. and Sellen, A. (to appear): "Making Contact in a Multi-media Environment", to appear in Behaviour and Information Technology.

Lövstrand, L. (1991): "Being Selectively Aware with the Khronika System", in *Proc. European Conference on Computer-Supported Cooperative Work ECSCW-91*, Amsterdam, Netherlands, October 1991.

Mackay, W. (1990): *"Users and Customisable Software: A Co-Adaptive Phenomenon"*, PhD Thesis, Sloan School of Management, MIT, Cambridge, Mass.

McGuffin, L. and Olson, G. (1992): "ShrEdit: A Shared Electronic Workspace", CSMIL Technical Report, Cognitive Science and Machine Intelligence Laboratory, University of Michigan.

Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T., and Wellman, B. (1991): **"**Experiences in the Use of a Media Space**"**, in *Proc. ACM Conf. Human Factors in Computer Systems CHI '91*, New Orleans, Louisiana.

Milligan, T. (1991): "The IIIF Design Report", University of Toronto, 1991.

Root, R. (1988): "Design of a Multi-Media Vehicle for Social Browsing", in *Proc. ACM Conf. Computer Support for Cooperative Work CSCW '88*, Portland, Oregon.

Stults, R. (1989): "The Experimental Use of Video to Support Design Activity", Xerox PARC Technical Report SSL-89-19, Palo Alto, Ca.