

Chapter 27

Reflecting on Several Metaphors of MUD-based Media Spaces

Saul Greenberg, Gregor McEwan and Michael Rounding

Abstract Over the last decade, we designed and used three media spaces: Teamrooms, Notification Collage, and Community Bar. All were oriented towards creating a shared environment supporting a small community of people: about two to around twenty members were expected to inhabit the media space. All provided others with a sense of presence through portrait images and/or snapshot-based video of its members, and all emphasised creation and sharing of real-time groupware artefacts. They differed in that each was designed around a different metaphor: multiple rooms for Teamrooms, a shared live bulletin board for the Notification Collage, and an expandable sidebar that contained multiple places for Community Bar. This chapter briefly reflects on how the systems and their metaphors served as a communal place. We saw that many factors – both large and small – profoundly affected how these media spaces were adopted by the community. We also saw that there was a tension between the explicit structures offered by media space design (rooms, places, bulletin boards and so on) vs. the very light weight and often implicit ways that people form and reform into groups and how they attend to information in the real world.

S. Greenberg
University of Calgary

G. McEwan
National ICT Australia, NICTA

M. Rounding
SMART Technologies, ULC

A version of this paper will appear in:

Greenberg, S., McEwan, G. and Rounding, M. Reflecting on Several Metaphors of MUD-based Media Spaces. In Harrison, S. (Ed.) Media Space: 20+ Years of Mediated Life. Series: Computer Supported Cooperative Work, Springer.

For information on this book, see:

<http://www.springer.com/computer/user+interfaces/book/978-1-84882-482-9>

1 Introduction

Media spaces come in many forms and flavours, but all try to encourage awareness leading to informal interaction between people (typically co-workers or collaborators) who benefit from casual interactions with one another. In the classic media space experiments – the late 1980s and early 1990s – such systems typically comprised always-on or easily available video / audio connections between interested parties (see Bly, Harrison, and Irwin 1993 for examples). Over time, these were augmented by computational tools that allowed people to move from conversation to actual work over computer artefacts. Around the same time, another popular class of social space was evolving: Multi-User Dungeons (MUDs and MOOs). Born of adventure games, traditional MUDs were text-based systems: game controllers created a space by using a server that hosted a variety of different rooms, each with a unique description and set of objects. A person could enter any number of different rooms, chat with other people in those rooms, and type commands to create and modify objects in the rooms. Multi-user virtual environments (MUVes) visualized MUDs as graphical worlds where people could present themselves as avatars that could navigate the world and encounter others (e.g., DIVE, Carlsson and Hagsand 1993). Text or audio connections to others were typically triggered by being collocated within a room or by proximity. MOOs combined MUDs and MUVes within a richer graphical user interface, and sometimes augmented their communication so that people could converse over richer channels. For example, the Jupiter project (Curtis and Nichols 1993) added MBONE audio and video conferencing, as well as graphics capabilities through shared whiteboards. More recently, Muds have been transformed yet again into the very popular MMORPG: massively multiplayer online role-playing games, with examples being World of Warcraft and Second Life.

Technically, there is not that much difference between a video-based media space augmented with computation tools vs. a MOO augmented with video; both serve the same purpose of providing a social world to its members by giving them awareness of who is around, and using that awareness to move into conversation and interaction. However, there are two key differences.

1. *Real vs Virtual Worlds*. Video-based media spaces are centered on the real world environment. The video connects two or more physical spaces so people appear somewhat co-located; add-on computational tools provide additional ‘virtual’ resources to the group. In contrast, MOOs are centered on a ‘virtual’ environment. People inhabit the virtual world, and add-on video / audio capabilities provide additional ‘real world’ connectivity.
2. *Intimate Collaborators vs Loosely Knit Communities*. Video-based media spaces tend to focus on very small groups of intimate collaborators: members are often goal-oriented and have a real need and desire to stay connected. In contrast, MOOs supported large, loosely-knit virtual communities, where anyone can enter the space (sometimes anonymously or with pseudonyms).

Our own interest was to merge the two approaches, where we wanted to re-fashion MOO-like virtual environments so they would better fit the real-world needs of modest-sized groups of intimate collaborators. The primary problem was not technical, for the capabilities of both video-based media spaces and MOOs began to overlap as they evolved. Instead, the challenge was how to redesign MOOs so they fit this different audience. Between about 1995 and 2007, I and my students developed, designed and used three MOO-based media spaces that offered both video snapshots and groupware artefacts. While each has superficially similar capabilities, they are designed around quite different metaphors.

- **Teamrooms**, commercialized as Teamwave Workplace, is based on the notion of multiple rooms (Greenberg and Roseman 2003; Roseman and Greenberg, 1997);
- **Notification Collage** is a shared live bulletin board viewable on a large public display and from people's workstations (Greenberg and Rounding 2001);
- **Community Bar** is an expandable sidebar that holds multiple places (McEwan and Greenberg 2005, Romero, McEwan and Greenberg, 2006).

This paper briefly reflects on each system – and each metaphor – as a communal place.

2 Teamrooms

Similar to MOOs, Teamrooms was designed around a rooms metaphor. However, our goal was to provide multiple virtual rooms that exploit features inherent in physical rooms used for team purposes (e.g., team rooms, war rooms, etc.). Its interface, features, and use are fully described in (Greenberg and Roseman 2003; Roseman and Greenberg, 1997). Figure 1 shows a screen snapshot of the commercial version of Teamrooms, called TeamWave Workplace. Some of Teamrooms' key ideas included:

- *a bounded space* that affords *partitioning* into a collection of rooms;
- *containment* through individual rooms, where they collect people and groupware objects;
- *permeability* of rooms allowing people and things to enter and leave them;
- *persistence* of objects within the room over time;
- *socially mediated ownership* that controls who should enter and use that room and how privacy is managed;
- *customization* of that room by how its occupants create and manipulate objects within it;

4 Reflecting on Several Metaphors of MUD-based Media Spaces

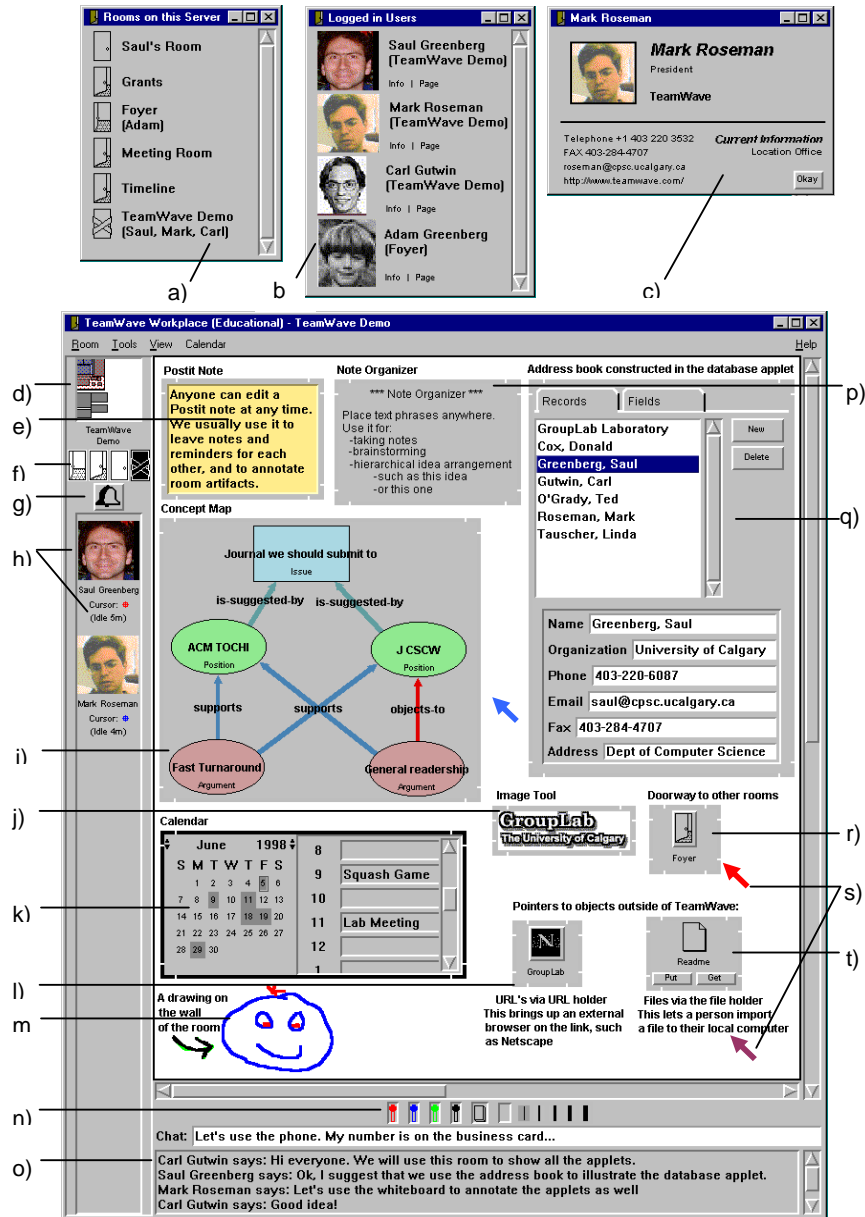


Fig. 1 TeamRooms / TeamWave Workplace user interface, showing a room and peripheral windows as seen by user Carl.

- *spatial location* where objects and people within a room are spatially positioned in a way that maintains common reference and orientation, and where proximity influences action and reciprocity;
- *habitation* where people can be aware of others across and within rooms, and where they can inhabit particular rooms.

Figure 1 illustrates the main components of its user interface. Figure 1a displays the “Rooms on this Server” window that lists all rooms currently available to the community, who is in it, and even a degree of privacy as suggested by the door icon. Figure 1b shows similar information, but in this case as a list of logged-on people identified by name, photo, and what room they are in. Clicking a person’s name reveals their business card (1c). A person enters a room by selecting that room from the list. The large window on the lower half of Figure 1 shows one of the many rooms created by this community. In this case, user Carl has entered a room called “TeamWave Demo”; Saul and Mark are also present in this room as shown by their icons (Figure 1h).

Rooms have many resources. The side bar includes a radar overview (1d), the ability to set privacy via the door state (1f), the ability to attract attention by ding-ing others (1g), and a list of people in the room, their cursor shape and idle time (1h). Depending on the system version, the images in 1h are static photos or video images updated every few moments in real time. The bottom bar includes a text-based chat dialog (1o) and a set of drawing implements (colored pens, eraser, and line thickness).

The center area is a groupware space: all see each other’s cursors in this space, all can work simultaneously, and all actions and artifact changes are seen by others in real time. The back wall is a sketchpad, and people can draw, erase and type on it (e.g., 1m). People can add a variety of special purpose applets to the room, such as a Postit™ note that serves as a multi-user text editor (1e), a groupware concept map editor (1i), a groupware calendar (1k), a note/list editor (1p), and/or a groupware database (1q). Within this space, they can also place and retrieve files (1t), URLs (1l), images (1j) and even doorways to other rooms (1r). Other tools (not shown) include a groupware web browser, a groupware file viewer, and even collaborative games. What is important is that the act of entering the room automatically connects these people together, where they can immediately see each other and all the things in the room, and where they can immediately chat and simultaneously interact over the groupware artifacts. Rooms and their artifacts persist, so people can come and go as they please.

We thought that groups would construct social places within these rooms, as the system no longer had many of the ‘seams’ found in conventional groupware. Rooms could serve as a place for both individual and group work; the distinction between the two was simply a matter of who occupied the room and the purposes the room was used for. Rooms also encouraged modeless interaction: real time interaction was just a consequence of people inhabiting the same room at the same time, while asynchronous interaction was a consequence of how people left arti-

facts (i.e., groupware objects showing content) within the room for others to see. Rooms would also let the social place develop over time; because things persisted (including writing that people could put on its back wall), people could craft the social meaning of the room by how they included objects within it, and how they decorated it. The collection of rooms would also form a community: while access control dictated who was allowed into a particular collection of rooms, any community member, once in Teamrooms, could create a room, could enter other rooms, and could see who was around. That is, access within a community was mediated through social vs. technical protocol.

In spite of the rich intellectual premises behind its design, Teamrooms did not live up to its promise as a social environment. While people did create their own rooms, we saw little actual interaction over time. Eventually, the commercial version of this product – Teamwave Workplace – was pitched as a place to hold planned classroom meetings rather than as a media space supporting casual social interactions and on-going work.

We believe that Teamroom's shortcomings were not with the room metaphor, but with the ways rooms were realized within it. With hindsight, we identified two major problems. The first was that Teamrooms did not effectively support awareness leading to casual interaction. A person could see who was around and thus available for interaction only *after* they actually logged into the system. Because logging in was relatively heavyweight, people would rarely do it just to see if someone was there. As well, people would not leave the system up and running just for awareness purposes, as it consumed considerable screen real estate. This defeated the 'always on' premise behind most media space designs. Thus there was little opportunity for casual interactions simply because no one was in a room long enough for others to notice. Another way to think of this is that TeamRooms was too MOO-like; people had to enter and inhabit the virtual world before they could see and interact with others. Unlike video-based media spaces, Teamrooms did not connect people's real world activities.

The second problem was that Teamrooms did not really support actual work. It only had 'toy' applications within it. While people could do simple tasks, they could not really share their real work done with commercial applications such as Microsoft Word, Excel, and so on. As well, voice was not supported, meaning that people would have to use an awkward chat system to mediate their real time interactions over these applications.

3 Notification Collage

To partially solve these two problems, we determined that our next system should somehow stress real-world (vs. virtual world) social interaction, and information sharing relevant to the group (vs. 'toy' groupware applications). First, we decided to base our design around the metaphor of a public bulletin board that would al-

ways be visible on a person's screen or on a public display; the idea was that this always-on visible bulletin board would exist within the person's real world environment and context. Second, the bulletin board would contain a collage of interactive information fragments, called *media items*, which are interactive groupware applications that let individuals post information they thought relevant to the group, where others could view and manipulate the media item's content. The result was a new groupware system called Notification Collage (Greenberg and Rounding 2001).

Notification Collage (NC) is illustrated in Figure 2 and works as follows. Distributed and co-located colleagues comprising a small community create a central server (a fairly trivial process). Each person in that community then connects to this server via an NC Client, which appears as a large window - a real-time collaborative surface - on their screen (Figure 2). Because of its size, we recommended people place NC on a second monitor located at their periphery. At this point, individuals can post media items (selected from the Posting Menu, Figure 2), and all members see these immediately. Akin to collages of information found on public bulletin boards, NC randomly places incoming elements onto this surface. While all see the same items, people can rearrange them as desired on their individual displays, e.g., in order to increase their visual prominence or hide less interesting ones from view. In particular, items placed on the right of a separator bar are never covered by new items. As illustrated in the figure, people can post assorted media: live video from desktop cameras; editable sticky notes; activity indicators; slide shows displaying a series of digital photos, snapshots of a person's digital desktop, and web page thumbnails. Some items allow people to move into direct interaction: people can move into an audio / video conference by selecting a person's video, and they can share a person's desktop by selecting a particular desktop image.

Unlike Teamrooms with its many virtual rooms as social places, this metaphor gives a group a single public place that holds meaning to them. First, it serves as a combined media-rich bulletin board, chat room, and video-based media space. We hoped that their focus on this single place would encourage sufficient postings and interactions to make it worth keeping always on, always visible on their display, and thus always present. Second, because it is a single bulletin board, we could post it in a large public display situated in a meaningful location, as well as on people's individual workstations, e.g., in a room or hallways populated by co-workers who are part of the NC community. Thus people could see its content as they walked by, or engage with others over it. Third, the overlap of items inherent in a large collage acknowledges that there may be a large number of information fragments, too many to tile neatly on the display. Finally, collages are customarily used to present unstructured information comprising diverse media, conceding that awareness information comes in many forms.

8 Reflecting on Several Metaphors of MUD-based Media Spaces

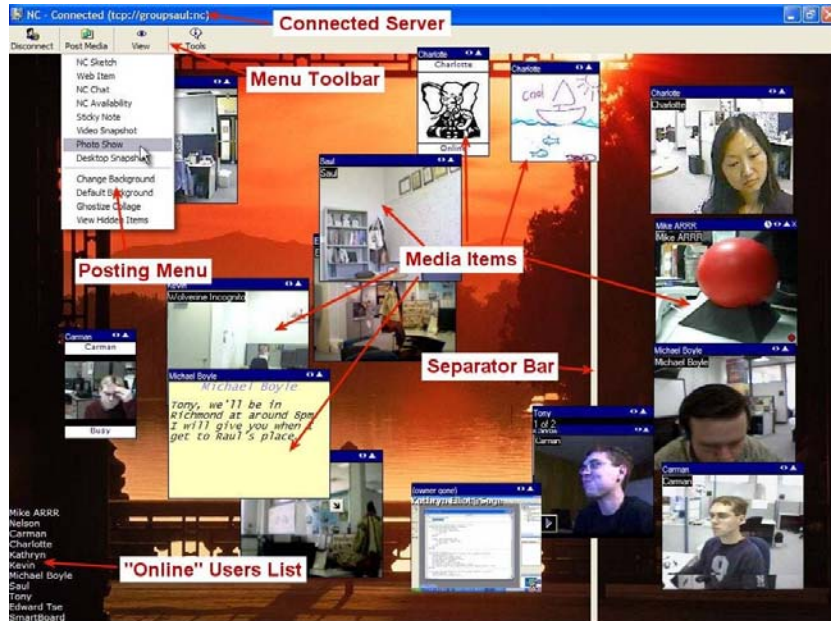


Fig. 2 Notification Collage

User experiences show that NC did evolve as a communal place, and that it served as a rich resource for awareness and collaboration. It gave people a keen sense of presence, especially because most community members chose to indicate their presence to others by posting live video. People's instinct was to create a visible presence for themselves: they wanted to see others, and others to see them. We also saw that media items triggered interaction. People acted on its information by engaging in text and video conversations. Unlike instant messaging and conventional media spaces, conversations sometimes began from people seeing interesting artifacts within the space and wanting to talk about them (e.g., photos or desktop snapshots). Next, the public nature of all actions encouraged interaction. All people could overhear conversations and see all postings; because even directed conversations and postings were visible to the group, anyone could monitor and join in. Furthermore, those cohabiting a public physical space could tell a collocated person about a note addressed to them. We also saw that media items concerning communication and information sharing (vs. the work-oriented groupware of Teamrooms) encouraged social engagement. People posted items they believed would interest others, such as desktop snapshots, announcements and vacation photos. Finally, the public display acted as a way for telecommuters to reach people (including room visitors) visible from its attached camera, and for those people to respond.

While successful as a place supporting a single small community, the Notification Collage had several limitations that restricted how it could be used by less well defined groups. As a single public place, it was all or nothing. People were either ‘in’ or ‘out’ of this community. This meant that people on the periphery of this group were sometimes reluctant to join in. The group was also very conscious of the appearance of ‘strangers’ (usually a friend of only one group member), where conversations would cease until that person was somehow introduced by an ‘in’ group member. Similarly, Notification Collage did not really support ad hoc groups. People were either a member of the community, or they were not. Its interface also proved somewhat heavy-weight in terms of how awareness was supported. People had to constantly review its contents to see what (if anything) had changed since they looked, although appearances of new items or changes to an item’s content typically stimulated a glance. Notification Collage was also more space-intensive than we would have liked. People without second monitors were severely disadvantaged, as the large size of the main window competed with other foreground applications.

4 Community Bar

The Community Bar (CB) (McEwan and Greenberg 2005; Romero, McEwan and Greenberg, 2006) extends our earlier work by trying to overcome the limitations of the Notification Collage while still building on the successes of media items. In particular its design is theory driven, where it is built around the Locales social science framework (Fitzpatrick 2003) and the Focus and Nimbus model of awareness (Rodden 1996). Generally, CB supports ad hoc groups by letting people create and enter *locales* (called *Places* in CB jargon). CB also balances providing rich awareness information *vs.* space requirements through its use of a *sidebar metaphor* that leverages the query in depth properties of the Microsoft Sideshow awareness display (Cadiz, Venolia, Jancke and Gupta 2002). Both these design considerations are described below.

The Locales Framework (Fitzpatrick 2003) suggests that people inhabit multiple social worlds, where each ‘world’ contains not only people, but offers a site and a means for their interactions. CB supports multiple locales through rapid creation of ‘Places’. For example, the particular individual’s CB client in Figure 3 displays two Places (i.e., two sites) called ‘mike test’ (top) and ‘CSCW class’ (bottom). Each comprises different sets of media items representing the various people who inhabit each place (e.g., the Presence media item in Figure 3 shows each person’s live video or image) and various means (e.g., people communicate through the Chat item; they share web pages through the Web item, they can post personal photos through the Photo item; they can even share their screens through the Screen sharing item). People can inhabit as many places as they wish. For example, Gregor, MB, and Mike R are in the ‘mike test’ place, while Kim, MB and KT are in the ‘CSCW Class’ place; MB cohabits both places. Long standing and

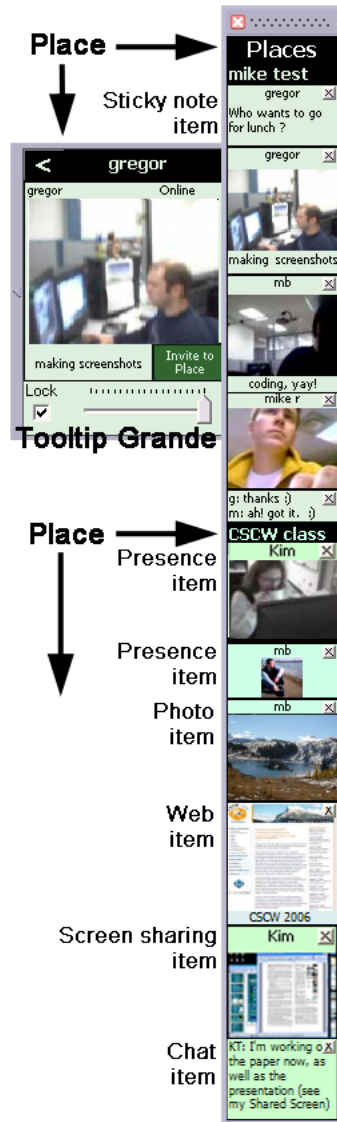


Fig. 3 Community Bar. Visible are two labeled places ('mike test' and 'CSCW class'), 5 types of items, and the Presence Tooltip Grande.

ad hoc groups can create, maintain, and destroy these places as needed. Through the media items, people within a place can present themselves to others, engage in conversation, and interact with group artefacts as desired. Each person can act in distinct ways in each of the Places they inhabit. Within a Place (and similar to MUDs), all information and interactions are public to all other people currently in that Place. Place members are able to share awareness information, to send broadcast queries (e.g. “Is there anyone who knows about X?”), and to overhear conversations and join those of interest to them. Unlike Notification Collage, CB supports multiple places rather than a single place. Unlike a room in Teamrooms, people can be in multiple places at the same time, and interact within any Place at leisure.

The Sidebar metaphor is important for lightweight transitions from peripheral awareness to foreground interaction. It recognizes the tension between a person’s desire for a *minimal* amount of unobtrusive yet dynamic awareness information of their intimate collaborators, against the need to act upon that information, e.g., to explore that information in depth, or to engage in rich communication as desired. Community Bar relieves this tension by offering people a progressive view of information. Rich yet not overwhelming awareness information is located at the periphery of the screen in a space conservative sidebar (shown in Figure 3, right). Moving the mouse pointer over items causes a “tooltip grande” to appear (Figure 3, left) that displays more information and provides interaction opportunities. Clicking on the tooltip grande title raises a “full view” permanent window (not shown in Figure 3) providing full information and interaction opportunities. To show this in more detail, Figure 4 shows this progression for the screen-sharing media item, which allows people to post part or all of their desktop to the group (Tee, Greenberg and Gutwin 2006a, 2006b). The tile view (left) shows a small image of Kim’s desktop, updated every several seconds. The size and resolution suffice to give others a broad sense of what she is doing. The Tooltip Grande (middle) offers a somewhat higher fidelity view of the same information at a faster update rate. The Full view (right) is at full fidelity; people can zoom into information and even request remote pointing to move into interactive screen sharing. Collectively, this progression of views allows the user to quickly stay aware of peripheral information, and to easily move into foreground interaction with information and people.

Finally, Community Bar uses the focus/nimbus model (Rodden 1996) to represent centre and periphery relationships. People express their involvement within a Place by using sliders to adjust both their nimbus (what others can see of them) and focus (how much they see of others). In this way, views and membership become somewhat more fluid. Unlike Teamrooms, where people are either in or out of a room, people can now adjust their focus/nimbus to control how much they are ‘inside’ a place.

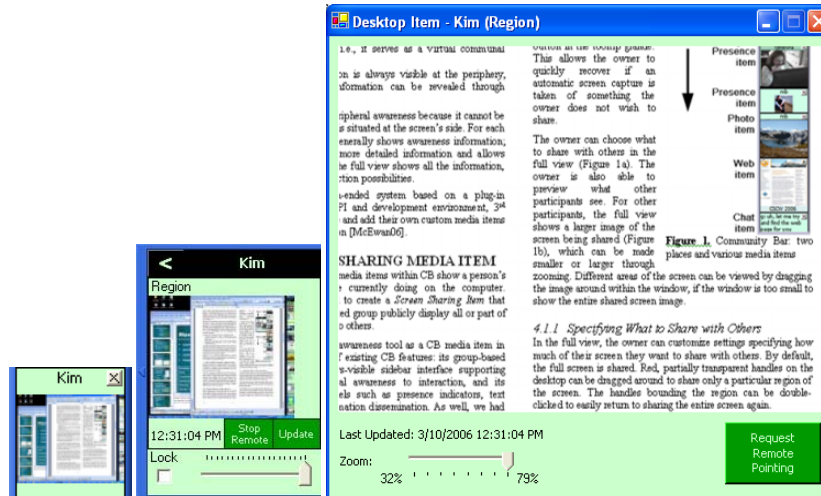


Fig. 4 The Tile, Tooltip Grande, and Full View of the screen-sharing item (Tee, Greenberg, and Gutwin 2006a, 2006b).

We performed a field study of CB in use (Romero, McEwan, and Greenberg 2006). Many things worked as predicted by the Locales theory, in particular, how people were able to maintain awareness and how they could move into interaction with others. However, the multiple Place functionality was not used heavily by this user community. We initially thought this was because the group was fairly cohesive, where they enjoyed working within one large Place (i.e., as in the Notification Collage). We thought this group did not see a strong need to splinter themselves into long-term sub-groups. Yet on closer inspection, we found that our study participants were easily divided into two groups: a ‘core’ group who often worked together closely and interacted with each, and a peripheral group comprising everyone else who had less work ties to the first group. This led to a divide in how CB was considered. Core group members consistently talked about the sense of belonging to the community that CB gave them. In contrast, peripheral members often reported that they felt like outsiders, and that most of the explicit communication on CB did not involve them directly. This was not necessarily a bad thing, as all people, whether core or peripheral, expressed sentiments on how useful CB was for maintaining an idea of what was going on with the rest of the group. We would have thought that this difference in member makeup would have encouraged multiple places. Yet when asked why they did not create new places, participants responded in very similar ways, saying that they were not needed in the existing community social structure. When asked about the situations under which they would use different places, most participants hypothesised that they would use different places if they were also involved in *distinctly different groups*

that did not know each other. That is, a CB Place seemed to define a community rather than a public locale.

A deeper analysis of CB use revealed that there actually were multiple locales in use within it, but this happened implicitly within a single Place. We realized that people tended to use sub-collections of media items as implicit locales, where they would ‘tune in’ to media items of interest and ‘tune-out’ items that were of lesser interest. They also formed implicit ad hoc groups as a function of their awareness and CB activities. For example, this was evident by the way chat items were used. Typically, only subgroups partook in discussions in chat items, and different chat items were often created (or taken over) for different purposes and people. Similarly, different sub-groups were interested in different things at different times: this likely led to some of the differences in how people interpreted some media item awareness information as useful vs. as clutter and distracting. Yet people seemed comfortable – even those who were ‘on the periphery’ – of doing all this ad hoc group formation within the context of the larger CB community vs. within the explicit structure of CB Places. On reflection, actual CB use is somewhat akin to how people inhabiting a common physical area selectively attend only some activities within it, and how they rapidly form and quickly reform ad-hoc sub-group clusters.

5 Discussion

All three systems were built around the notion of a collection of public media items that portray people (usually as live video snapshots), their interactions (usually as text chats), and their things (usually as information containers or mini-groupware applications). They differ considerably in the metaphors they followed, which in turn affected how each system structured and presented these items. What we saw is that many factors – both large and small – profoundly affect how these media spaces are adopted by the community. In spite of its rich room metaphor, Teamrooms was not adopted as we had expected, simply because it lacked the lightweight awareness critical to casual interaction and because the ‘walls’ comprising its room were too hard – they isolated community members rather than brought them together. This left it more suitable as a meeting tool rather than an always-on media space. Notification Collage did work as it offered a rich multimedia space for casual interaction. However, it was limited as being an ‘all or none’ system; people were either in the group defined by the single media space, or out of it. Community Bar achieved the same effect as the Notification Collage while doing a better job at balancing awareness and distraction. Yet its key property – that of Places – was not used in the way we expected, i.e., Places defined the entire community vs. ad hoc groups. Still, we did see sub-groups evolve within a single Place through how people used its items and how they attended them.

5.1 Implicit vs. Explicit Structures

It seems there is a tension between the explicit structures offered by media space design (rooms, places, bulletin boards and so on) vs. the very light weight and often implicit ways that people form and reform into groups. In real life, we do this through a variety of social mechanisms - by physically moving closer together, by sharing things in particular ways, by cohabiting a space, by moving between multiple spaces, and by selectively attending and responding to the world around us. In the computer world, these everyday physics don't apply. Instead, we substitute explicit structure through our metaphors as a mitigating surrogate that anticipates how groups form and reform, and that controls the social mechanisms of groups. In practice, we see that these explicit structures are often ignored or become hurdles. In contrast, systems with little structure do seem to work better because people use their own attentive and social resources to define their group in a subtle, highly flexible, and tacit way. Yet we expect an unstructured approach will have problems, as they likely will not scale beyond small reasonably cohesive groups. A challenge is how we can provide explicit structure as a flexible backbone to peoples' social activities, where it suggests rather than forces people to interact in a way that works well for them.

5.2 Transitions between Real vs. Virtual worlds

Let us now reconsider the main difference between media spaces and MUDs. As mentioned, media spaces connect intimate collaborators and are centered around the real world; the idea is that because people live in the real world, technology can create a virtual portal connecting geographically separated real world locations. In contrast, MUDs collect loosely knit communities centered around a virtual world; the idea being that this virtual place transcends the physical boundaries and physics that could otherwise limit community formation and interaction.

While both views appear reasonable, they are incomplete in of themselves, as they assume that people primarily 'live' in either the physical world (which favours the media space approach) or the virtual world (which favours the MUD approach). Yet the reality today is that most people live in both places. What has happened since the late 1980s (when media spaces and MUDs were first introduced) is that the desktop computer has become an integral part of people's physical world. Many people are now heavy computer users, spending at least a few hours a day at their screens, or referring to them repeatedly over the course of a day. Yet these same people inhabit a physical world, where they also attend to life off the screen. Thus computer use is interwoven with physical activity.

On reflection, we realize now that the evolution of our three systems moved their design towards a space that somewhat blends the distinctions between media spaces and MUDs, as they were progressively centered around having a real world component that transitioned into the virtual world. This was a good thing. For example, CB and Notification Collage were successful as awareness systems as their always-on state meant they 'lived' in the real world. Its information was always available, where people habitually monitored their display while pursuing some

other real-world activity or while working at their computer. The light-weight manner that people could attend this information within the Notification Collage and CB place better reflected their real-world engagement with information and with the other people inhabiting the place. CB's sidebar also did a better job of managing people's real world awareness needs, where they could move from information easily available in the real world (i.e., via a quick glance), to more in-depth information held by the virtual world (e.g., by navigating to tool-tip grandes and full interaction views).

The failures we noted previously show where transitions between the real and virtual world were not well supported. First, the state of activity in a Teamrooms room was visible only if one was in the place; in practice, this meant that critical information was rarely accessible 'at a glance' from the real world. Out of sight became out of mind. Second, both the Notification Collage and Teamrooms demanded a large amount of real estate. People mediated this by either devoting a large part of a second display to it (which was expensive in terms of cost and desk space), or by covering up the window with other windows or shrinking it down to the toolbar if the screen was used for other things (which fundamentally changed the ease of accessing its information). When the window was in full display, it was visually dominant, i.e., it demanded attention in a way that could have been out of proportion to how people wanted to use it. If people expended excessive time monitoring activity in the virtual world that did not match its real world importance, then it would become a distraction. The result is that people would likely turn off the system or shrink it out of view or simply not use it. Third, the explicit structure of a virtual group suggested by both a Teamrooms' room or a CB place was largely ignored, as it often did not reflect how people perceived, managed and attended their real-world groups.

5.3 Embodied Interaction

Dourish's (2001) theory of *embodied interaction* can help us reconsider the phenomena noted above. Embodied interaction is the way that physical and social phenomena unfold in real time and real space as a part of the world in which people are situated, right alongside and around us, and how people create meaning about the world through their own actions and the actions of others.

Under this theory, we should now understand that people are including their virtual systems as part of these physical and social phenomena. We suspect that the specific successes as seen in NC and CB have more to do with having some portion of these system – usually the attentional devices that promote awareness – embedded in the real world where people were doing all of their other activities. People's world as a whole thus incorporated a complex mixture of social, physical and virtual components. Each system provided an extension from the physical to the virtual world in which people were currently acting, with smooth transitions in and out of this new space. In contrast, MUDs and Teamrooms created a separate world to inhabit, where the transitions are heavyweight. Thus people had to choose to be embodied in one world or the other. While appropriate for dedicated activities (such as meetings), this split embodiment does not work well for casual

interaction. Again we come back to and stress the need for transitioning from informal awareness to casual and formal interaction if the system is to become a social place. It is not a matter of connecting our physical worlds (as done in media spaces) *or* connecting our social worlds (as done in MUDs), but rather of blending those worlds together in a way that eases the transitions between them.

While our three systems all worked to varying extents, none are optimal. The design challenge is how to strike a balance between the offerings of the virtual world and how they manifest themselves in the embodied interaction of the moment-by-moment real world. Perhaps ubiquitous computing approaches may help this process, where we may see virtual information and people situated in our real world in a way that preserves our natural social ways of interacting with the world around us (Dourish 2001; see Greenberg and Kuzuoka for an example).

Acknowledgments

Research was partially funded over the years by the NSERC Discovery Grant, the NSERC NECTAR Research Network, by Microsoft Research, and by the NSERC/iCORE/Smart Technologies Chair in Interactive Technologies.

Software Availability

Community Bar and the Notification Collage are available for download at <http://www.cpsc.ualgary.ca/grouplab/cookbook/>. Both were designed around earlier versions of Microsoft Windows.

References

- 1 Bly S, Harrison S, and Irwin S (1993) Media spaces: bringing people together in a video, audio, and computing environment. *Communications of the ACM*, Volume 36 Issue 1, January
- 2 Cadiz JJ, Venolia GD, Jancke G, and Gupta A (2002) Designing and deploying an information awareness interface. *Proc ACM CSCW (2002)* 314-323
- 3 Carlsson C and Hagsand O (1993) DIVE - A platform for multi-user virtual environments. *Computers and Graphics*, 17(6). 663-669.
- 4 Curtis P and Nichols D (1993) MUDs grow up: Social virtual reality in the real world. In *Proceedings of the Third International Conference on Cyber-space*
- 5 Dourish P (2001) *Where the Action is: The Foundations of Embodied Interaction*. MIT Press, October
- 6 Fitzpatrick G (2003) *The Locales Framework: Understanding and Designing for Wicked Problems*. Kluwer Academic Publishers
- 7 Greenberg S and Kuzuoka H (2000) Using Digital but Physical Surrogates to Mediate Awareness, Communication and Privacy in Media Spaces. *Personal Technologies*, 4 (1), January, Elsevier

- 8 Greenberg S and Roseman M (2003) Using a Room Metaphor to Ease Transitions in Groupware. In M. Ackerman, V. Pipek, V. Wulf (Eds) *Sharing Expertise: Beyond Knowledge Management*, 203-256, January, Cambridge, MA, MIT Press
- 9 Greenberg S and Rounding M (2001) The Notification Collage: Posting Information to Public and Personal Displays. *Proc ACM CHI*, 515-521, ACM Press
- 10 McEwan G and Greenberg S (2005) Supporting Social Worlds with the Community Bar. *Proc ACM Group 2005*, ACM Press
- 11 Rodden T (1996) Populating the Application: A Model of Awareness for Cooperative Applications. *Proc. ACM CHI* 88-96
- 12 Romero N, McEwan G and Greenberg S (2006) A Field Study of Community Bar: (Mis)-matches between Theory and Practice. Report 2006-826-19, Department Computer Science, University of Calgary, Calgary, Alberta, Canada, T2N 1N4, March 17th
- 13 Roseman M and Greenberg S (1997) A Tour of TeamRooms. Video Proc ACM CHI, ACM Press. Videotape
- 14 Tee K, Greenberg S, and Gutwin C (2006a) Providing Artifact Awareness to a Distributed Group through Screen Sharing. *Proceeding of the ACM CSCW'06 Conference on Computer Supported Cooperative Work*, ACM Press, November
- 15 Tee K, Greenberg S, Gutwin C, and McEwan G (2006b) Shared Desktop Media Item: The Video. *Video Proceedings of ACM CSCW'06 Conference on Computer Supported Cooperative Work*, November, ACM Press. Video and two-page summary. Duration 4:00

About the authors

Saul Greenberg

Professor and iCore Chair
Human-Computer Interaction & Computer-Supported Cooperative Work
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
saul.greenberg@ucalgary.ca

Saul Greenberg is a Full Professor in the Department of Computer Science at the University of Calgary. While he is a computer scientist by training, the work by Saul and his talented students typifies the cross-discipline aspects of Human-Computer Interaction, Computer-Supported Cooperative Work, and Ubiquitous

Computing. He and his crew are well known for their development of toolkits enabling rapid prototyping of groupware and ubiquitous appliances, innovative and seminal system designs based on observations of social phenomenon, articulation of design-oriented social science theories, and refinement of evaluation methods. His research is well recognized. He holds the iCORE/NSERC/Smart Technologies Industrial Chair in Interactive Technologies. He also holds a University Professorship, which is a distinguished University of Calgary Award recognizing research excellence. He received the CHCCS Achievement Award in May 2007 and was also elected to the ACM CHI Academy in April 2005 for his overall contributions to the field of Human-Computer Interaction. Saul is a prolific author who has authored and edited several books and published many refereed articles, as listed at <http://grouplab.cpsc.ucalgary.ca>. He is also known for his strong commitment in making his tools, systems, and educational material readily available to other researchers and educators.

Gregor McEwan

National ICT Australia, NICTA

Gregor McEwan is currently a Research Engineer at National ICT Australia (NICTA), with a full-time commitment to the (bracetto) project. Before starting with NICTA in September 2006, he was completing his M.Sc. at the University of Calgary in Canada. Gregor's expertise is in Computer-Supported Cooperative Work with a special emphasis on creating groupware environments for small communities. His recent graduate work applies social theories of interaction to groupware design. Previous to his graduate work he was involved in designing and implementing a complex groupware environment, and also research into social communities formed around sharing of web resources. At that time, he was also involved in research about combining Participatory Design and eXtreme Programming development methodologies. Gregor has previously worked as a Research Scientist at the Distributed Systems Technology Centre (DSTC) and was involved in the Information Ecology, Ambience, Elvin and Orbit projects.

Michael Rounding

SMART Technologies, ULC

Michael Rounding is currently a User Experience Specialist with SMART Technologies where he researches, designs and evaluates interactive technologies used globally in classrooms and corporate boardrooms. He is an active field researcher, and teaches courses internal to SMART regularly. He received his MSc in 2004 at the University of Calgary in Computer Science, specializing in Human Computer Interaction, working with Dr. Saul Greenberg on groupware interfaces and architectures used to promote communication and collaboration in groups of partially co-located and distance-separated collaborators. In addition to his current work with SMART, he has also applied his training to the Medical Imaging industry, designing imaging products to help doctors improve patient care.