Chapter 13
Mediaspace – Meaningspace – Meetingspace[1][2]

Bill Buxton

Thoughts exchanged by one and another are
not the same in one room as in another.
Louis I. Kahn

Abstract As technology becomes ever more pervasive in our lives, one of the fundamental questions confronting us is how to resolve the increasing complexity that too often accompanies it – complexity which threatens to prevent our reaping the potential benefits that it offers. In addressing this question, much of the literature has focused on improving the design and usability of the interface to the technologies themselves. In this paper we investigate another approach, one in which some of the complexity in using the devices is eliminated by exploiting some of the key properties of architectural and social space. Our work is based on the observation that there is meaning in space and in distance. Hence, we can relieve users of the complexity having to explicitly specify such meaning, since – through appropriate design – it can be implicit, given its spatial context.

Introduction

When you walk into a lecture hall at a university, even one that you have never been in before, and where you know nobody, you still know who is the professor and who are the students. If you see a photo of a dinner party, with everyone sitting around the dining room table, you know who are the hosts are and who are the guests. Walking in the park, you can tell if two people are in love, even if you see them only from a distance.

In each of these examples, we know what we know because of our literacy in the meaning of space. In the lecture hall, the professor is at the front, and the students in the chairs. We gain our understanding from the position of the people relative to the architectural space. With the dinner party, we can infer who are the hosts because they typically sit at the head of the table. In this case, it is position relative to a fixed object in the architectural space that provides the cues for interpreting the social relationship amongst the party. And finally, with the lovers in the park, it is their physical proximity relative to each other – regardless of if they are in the park, on a bus or on a boat – which leads to our conclusion about their emotional closeness.

What all of these examples illustrate is that from a lifetime of living in the everyday world, we have all built up a phenomenal depth of knowledge about the conventions of space and its meaning – both absolute and relative, and physical and social. This is knowledge that we exploit every day, in almost everything that we do, in order to make sense of, and function in, the world.

---

[1] Harrison, Steve (Ed.)(2009). *Media Space: 20+ Years of Mediated Life.* New York: Springer Verlag, 217 – 231.

[2] I have fixed a few minor typos and awkward wording in this version, but not changed the substance of the publisher version.

It is also something that can be exploited to reduce the complexity and intrusiveness of the technologies that we introduce into our world. This is something that the examples discussed in this chapter are intended to illustrate.

The examples discussed have been implemented and used in practice. The approach was opportunistic: to do smart things with stupid technologies. Rather than make engineering breakthroughs, our objective was to create an opportunity to gain experience living with these technologies *before* they were commercially viable. Our hope was that the human insights gained might help inform future design practice and development. Our mantra, while doing this work was as follows:

The only way to engineer the future tomorrow is to have lived in it yesterday.

Background

In the 1980s I was involved in two projects at Xerox PARC. One was the *Ubiquitous Computing* project led by Mark Weiser, which was to have a major impact on our thinking about the future of computation (Weiser 1991). The other was the *Mediaspace Project*, initiated by Bob Stults, Steve Harrison and Sara Bly (Stults 1986; Bly, Harrison & Irwin 1993).

The former had to do with digital computers, and as manifest at PARC at the time, primarily pen-based computing on three scales: palm-sized "tabs", slate-sized tablets, and whiteboard sized panels. All were networked using (then) uncommon wireless technologies (infrared and packet radio), and had high levels of interoperability.

On the other hand, the Mediaspace work had to do with audio/video technologies that let designers, in particular, better collaborate at a distance. The idea was to use the technology to establish a persistent sense of presence amongst a community that was geographically distributed. The technologies used were decidedly "old school" in that conventional analogue video gear (albeit controlled by a novel computer interface) formed the foundation of the system.

Despite both existing at PARC, these two projects were very far apart, physically and intellectually. Yet, in my mind, the two were actually two sides of the same coin. Both dealt with technologies that were destined to become pervasive. At the meta-level, the only difference was that the slant of one was computation and the other remote collaboration.

Between 1987 and 1989 I had the opportunity to design the media infrastructure for the new EuroPARC facility in Cambridge, what became known as the "IIIF" or "RAVE" system (Buxton & Moran 1990; Gaver *et al.* 1992). This gave me the chance to take an initial step in integrating some of these concepts. Then, from 1989 – 1994, I got a chance to go through another iteration when I set up the Ontario Telepresence Project in Toronto.

It is work undertaken as part of this latter project that forms the basis for this chapter. However, it is important for me to provide the above historical context since it is hard to separate what we did in Toronto from what was being done at PARC. In many ways they *were* the same project, since while I was scientific director of the Telepresence Project in Toronto, I was also working half time at PARC as part of both the UbiComp and Mediaspace projects. Furthermore, the software which provided the foundation for the Telepresence project was first developed at EuroPARC, then further developed in Toronto, and subsequently installed at PARC.

The Social-Spatial Anatomy of My Workspace

Let's start with my old office at the University of Toronto, which is shown schematically in **Error! Reference source not found.**. Even within this relatively simple space, very different social interactions or protocols are associated with each of the various locations identified in the figure:

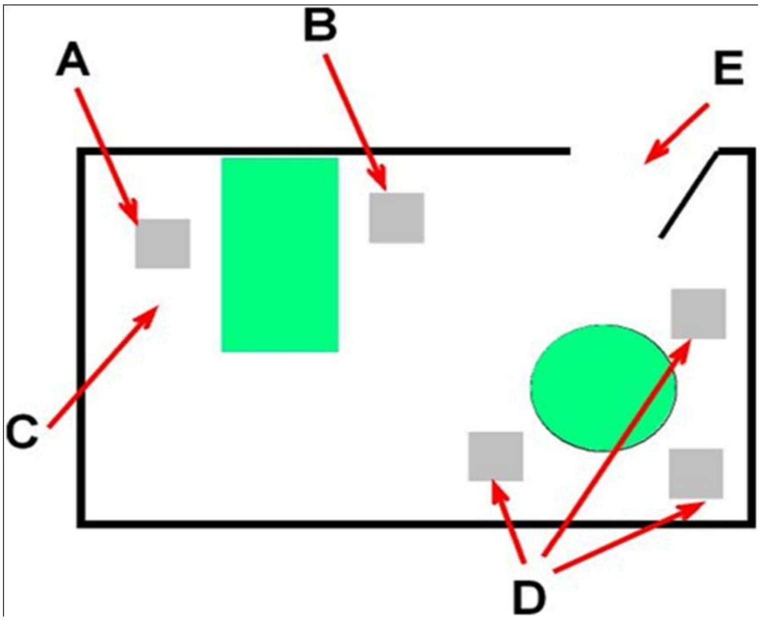

Fig. 13.1 Schematic of my office.A number of distinct locations in the office are indicated, including the chair behind my desk (A), the chair across from my desk (B), standing space behind my desk (C), and chairs around the coffee table (D) and the door (E). Different social functions are associated with each location. The deployment of any technology in the office to support collaboration or social activities should reflect and respect these differences.

A. My chair behind the desk.
B. The chair across from my desk.
C. The position beside my chair.
D. The chairs around the coffee table.
E. The doorway.

To get a taste of what I mean, consider a meeting with a student.

First, I might sit in my chair (A) and have the student sit across the desk from me in position (B). In this case, I am Professor Buxton, and they are not. This would likely be the situation if I were telling a student that they had failed, or if I was formally congratulating them on a great job.

Second, if I was working closely with the student on something, they might come behind my desk to position (C), while I sat in my chair. However, it would be very unusual for a stranger or someone with whom I was not working closely, or did not know, to go behind my desk.

Third, if I was having a casual meeting, or just chatting, we may sit around the coffee table in the chairs labelled (D). This would occur if the meeting was informal, and it would indicate that the relationship was more collegial than subordinate. It would be a meeting with "Bill" rather than "Professor Buxton."

Fourth, I may be working at my desk, and the student pop their head in the door to ask something. If I do not ask them in, the student would know that I was busy, not be offended, and that the conversation would be brief.

Finally, when the meeting involved a number of students, for example, rather than use my office, the meeting would take place in a conference room – a space which has its own set of conventions around space.

Our premise is that any technology introduced into such spaces must reflect and respect these space-function-distance relationships. Therefore, the appropriate technologies need to be distributed at the appropriate locations within that space.

Mies van der Rohe notwithstanding, the implication of this is, *More is Less*: interaction with *more* of the right technologies spatially deployed in the appropriate locations is much *less* intrusive that channelling everything through a single general-purpose technology typically anchored to a single, and therefore generally wrong, location. In designing from this perspective, we evolved a few basic principles, including the following:

Design Principle 1: *Maintain a clear distinction between person space and task space*

Design Principle 2: *Respect the function-location relationships and conventions for all present, either physically or via telepresence.*

Design Principle 3: *Treat electronic and physical "presences" or visitors the same.*

Design Principle 4: *Use the same social protocols for electronic and physical social interactions.*

We will now work through some examples that illustrate how we approached supporting scenarios such as those discussed above for those who were not physically present.

Example: At the Desk

Let's start with the scenario of working closely with someone at the desk. The typical configuration here is for the remote person to appear on a monitor by the desk that has a video camera placed on top of, or beside it. In our implementation, illustrated in Fig. 13.2, we already see Principle 1 kicking in, resulting from a departure from the norm. The monitor used is physically distinct from the computer screen on which the work being discussed would appear. Among other things, this prevents any contention for screen real-estate on the computer monitor. Documents and people have their own distinct place – just as in the physical world (Buxton 1992).

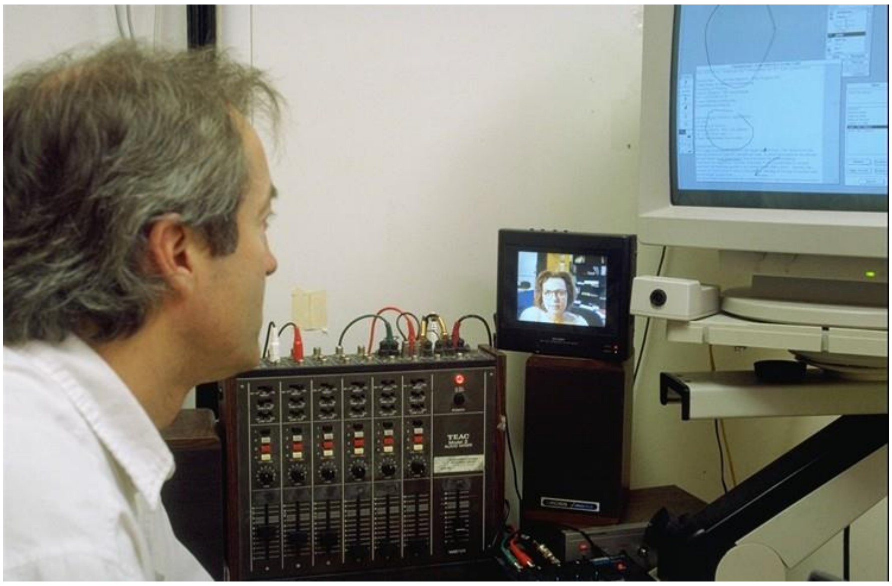

Fig. 13.2 A Typical Desktop Video Conferencing Configuration. Conferencing is typically channelled through a video camera on top of a monitor on the user's desktop. However, that monitor is distinct from the computer monitor so as to differentiate "person space" from "task space".

Apart from the problems of approach, which we will discuss later, this configuration could be fine for some meetings, such as those where it would be appropriate for the visitor to be positioned at location B or C in my office. But what if I am sitting at my coffee table, position D, a not uncommon thing? How do I interact with the remote person if the monitor on which they appear and the camera which they see me by is by my desk pointing at my chair at location (A)? I could get up and reposition either myself or the video set-up, but that would be a disruption that need not happen.

Example: Around the Coffee Table

Design Principle 2 leads us to our solution. Just as there are different places where those physically present can sit for different purposes, so should it be for the remote participant. Hence, besides the video system at my desk (Fig), there was also a system at the coffee table (Fig. 13.1) where a visitor could "sit" and participate in around-the-table conversations.

Fig. 13.1 Remote Participation in an Informal Group. Here a group, including a remote participant (detail in inset), are sitting around the coffee table in my office having a casual meeting. (In position "D" relative to the schematic in Fig. 13.1).

The visitor is able to sit at the physical location appropriate for the social function of the meeting, regardless of whether they are there physically or electronically, thereby supporting Principle 3.

Example: Approach and the View from the Door

So far so good. But there remains the small matter of how you entered my office in the first place. Social conventions are as much about transitions, such as approach and departure, as they are about being here or there. In the desktop video situation illustrated in Fig, for example, you are either there or not. When you are there, you are right in my face. Worse than that, when you arrive, you do so abruptly, in a way that violates normal social conventions of approach.

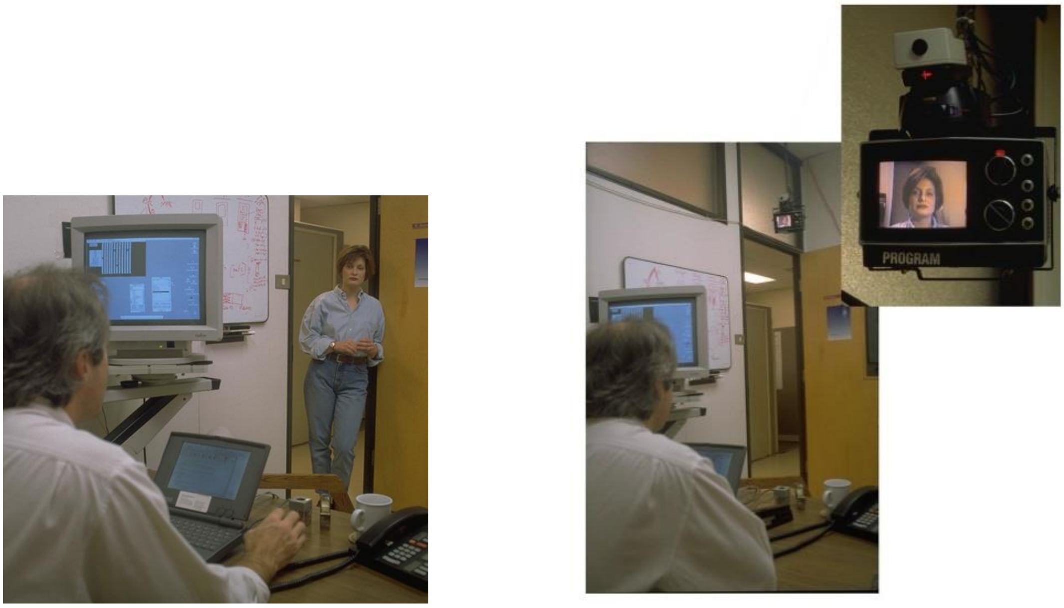

Fig. 13.2 Maintaining Social Distance. In establishing contact, one appears by the door and has a from-the-door view via the camera, regardless of whether one approaches from the physical corridor (left image) or the electronic corridor (right image). People approaching electronically do so via a monitor and speaker mounted above the door (inset on right image). The social graces of approach are preserved, and the same social conventions are used for both physical and electronic visitors.

Fig. 13.2 illustrates our approach to addressing this problem. When you come to my office, you come via location (E), the door. If you come physically, then all is normal. If you come electronically, you also appear by the door, but on a small video monitor mounted above it. In a manner analogous to hearing your footsteps coming down the physical corridor outside my office, I hear your approach *via* an emitted "earcon" which emanates from a speaker by the door monitor. And, I hear you *before* you appear or can see me. When you do see me (which is at the same time I can see you), you do so from a wide-angle low-resolution camera that is integrated with the monitor that I see you on and the speaker from which I hear you. Thus, the glance that you first get is essentially the same as what you would get through the door. If I am concentrating on something or someone else, I may not see you or pay attention to you, just as would be the case if you were walking by in the hall (even though I may well hear that someone is there or has passed by). Appropriate distance is maintained. If you knock or announce yourself, I may invite you in, in which case you could take a place at my desk, or around the coffee table, whichever is more appropriate.

Example: Front-to-Back-to-Front Videoconferencing

As stated previously, not all meetings happen in my office. So let us now look at how some of these ideas apply in a small conference room. Most such rooms equipped for videoconferencing that I have used are set up more-or-less like the one shown in Fig. 13.3.

Fig. 13.3 Front-to-Back Videoconferencing. Here the remote person is located at the front of the room. Hence, one could reasonably infer that he is presenting, since the front is the place that one presents from.

The videoconferencing technology is at the front of the room, which is the location from which one typically presents. The problem arises when the role of remote person appearing on that front monitor is one of attendee, rather than speaker – especially if there is a physically present speaker at the front. The remote attendee likely just sees is the back of the head of the presenter, and is simultaneously a visual distraction for the other attendees. Things are socially broken in this scenario because the remote person appears at a location which violates the moral order of the meeting.

Fig. 13.4 Back-to-Front Videoconferencing. Remote attendees to a meeting take their place at the table by means of video monitors mounted on the back wall. They see through the adjacent camera, hear via a microphone, and speak through their monitor's loudspeaker. The presenter uses the same conventional skills in interacting with those attending physically and those attending electronically. No new skills are required.

The wrong way to fix this, however, is to move the conferencing gear to the back of the room. This just reverses the problem. Alternatively, what we did was provide multiple "seats" at different locations. Thus remote participants could "take a place at the table" like any other participant when not presenting. An implementation of this "back-to-front" videoconferencing is illustrated in Fig. 13.4.

Here, a presentation is being made to one remote and four local participants. Due to the maintenance of audio and video reciprocity coupled with the appropriate use of space, the presenter uses the same social mechanisms in interacting with both local and remote attendees. Stated another way, even if the presenter has no experience with videoconferencing or technology, there is no new "user interface" to learn. If someone raises their hand, it is clear they want to ask a question. If someone looks confused, a point can be clarified – no matter where they are.

Example: Hydra: supporting a 4-way round-table meeting

Underlying our work was a kind of mantra that reflects the attitude of ubiquitous computing: one size doesn't fit all. We experimented with different designs for different types of meetings. One of these was a technique to support a four-way meeting, where each of the participants is in a different location. It was designed to capture many of the spatial cues of gaze, head turning, gaze awareness (Ishii, Kobayashi & Grudin, 1992) and turn taking that are found in face-to-face meetings. Consistent with the design principles outlined above, we do this by preserving the spatial relationships "around the table"[3]. This is illustrated in Fig. 13.5.

.

Fig. 13.5 Using Video "Surrogates". The photo on the left shows a 4-way video conference where each of the three remote participants attends via a video "surrogate." By preserving the "round-table" relationships illustrated schematically on the right, conversational acts found in face-to-face meetings, such as gaze awareness, head turning, etc. are preserved.

Each of the three remote participants are represented by a small video surrogate. These are the small Hydra units seen on the desk (Sellen, Buxton & Arnott, 1992; Buxton, Sellen & Sheasby, 1997). Each provides a unique view of one of the remote participants, and provides each remote participant a unique view of you. The spatial relationship of the participants is illustrated by the "round-table" on the right. Hence, relative to you, person A, B and C appear on the Hydra units to your left, front and right, respectively. Likewise, person A sees you to their right, and sees person B to their left.

Collectively, the units shown in the figure mean that the user has three monitors, cameras and speakers on their desk. Yet, the combined footprint is less than that of a conventional telephone.

[3] This idea of using video surrogates in this way for multiparty meetings turns out not to be new. After implementing it ourselves, we found that it had been proposed by Fields (1983).

These Hydra units represent a good example of transparency through ubiquity. This is because each provides a distinct location for the source of each remote participant's voice. As a result, due to the resulting "cocktail party effect", the basis for supporting parallel conversations is provided. This showed up in a formal study that compared various technologies for supporting multiparty meetings (Sellen, 1992). The Hydra units were the only technology tested that exhibited the parallel conversations seen in face-to-face meetings.

The units lend themselves to incorporating proximity sensors that would enable aside comments to be made in the same way as face-to-face meetings: by leaning towards the person to whom the aside is being directed. Because of the gaze awareness that the units provide, the regular checks and balances of face-to-face meetings would be preserved, since all participants would be aware that the aside was being made, between whom, and for how long.

None of these every-day speech acts are supported by conventional designs, yet in this instantiation, they come without requiring any substantially new skills. Again, there is no "user interface."

Finally, we can augment the basic Hydra units by placing a large format display behind them. As shown in Fig. 13.6 this is used to function like a large electronic "whiteboard" which enables the user to easily direct their glance among the other three participants and the work being discussed. Furthermore, if all four participants have their environments configured the same way, and the same information is displayed on each of the large displays, then each has optimal sight lines to the "whiteboard." Here is a case where the combination of electronic and physical space (Buxton, 1992) provides something that is an improvement on the traditional physical world where, if the physical whiteboard were across from you, it would be behind person "B" sitting opposite you. Furthermore, note that the awareness that each participant has of who is looking at who (so-called "gaze awareness") extends to the "whiteboard".

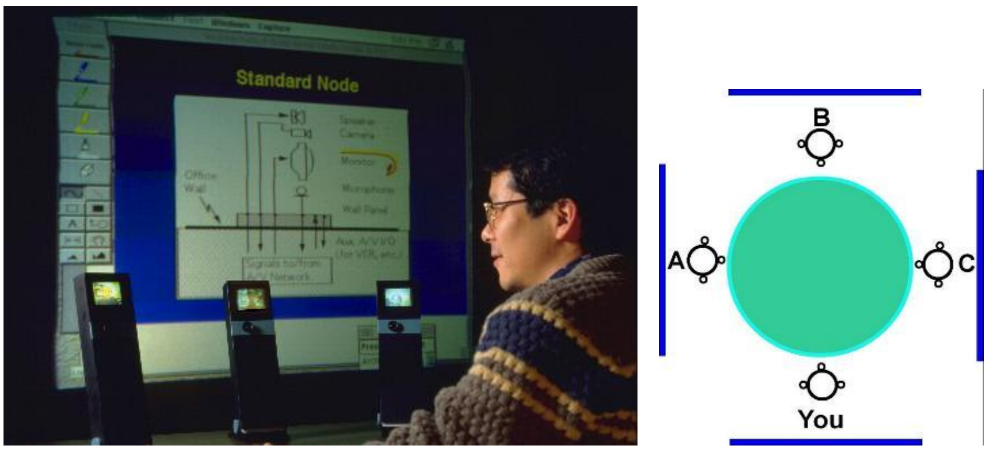

Fig. 13.6 Seamless Integration of Person and Task Space. The photo on the left also shows a 4-way video conference using the Hydra units. However, this time, a large electronic "whiteboard" containing the information being discussed appears behind the units. As illustrated in blue in the schematic on the right, the same display can appear behind the units at each of the four sites, thereby giving each participant ideal sight lines to the "same" whiteboard (something that does not occur in same-place round-table meetings.) Furthermore, gaze awareness now extends to whether one is looking up at the "whiteboard" or at a person, thereby seamlessly blending person and task space.

Example: Size Matters

Scale, as well as location is important in terms of its ability to affect the quality of interaction in a Mediaspace. Consider the impact of electronically sitting across the desk from one another, as illustrated in Fig compared to Fig. 13.7, where, the remote participant appears life-size across the desk. In this case, we are using essentially the same configuration as we saw in Fig. 13.6; however, in this case the large display is showing the image of the remote person in a 1-on-1 conversation. I am captured by the Hydra camera, but the large display replaces the Hydra monitor. A number of significant points arise from this example.

Fig. 13.7 Face-to-Face. In this scenario, each participant has a computerized desktop on which the same information is displayed. The intention is to capture the essence of working across the desk from one-another. Each sees the remote participant life-size. The video camera (from a Hydra unit) is unobtrusive on the desk. Participants interact with the computer using a stylus. When one participant looks down to their desktop, their eyes seem to project into the space of the other, thereby strengthening the sense of telepresence. While there is a considerable amount of technology involved, it is integrated into the architectural ecology. What one gets is lots of service and lots of space, not lots of gear and appliances.

First, it is not like watching TV. Due to the scale of the image, the borders of the screen are out of my main cone of vision. The remote person is defined by the periphery of their silhouette, not by the bezel of a monitor. Second, by being life size, there is a balance in the weight or power exercised by each participant. Third, and perhaps most important, the gaze of the remote participant can traverse into my own physical space. When the he looks down on his desk, my eyes are directed to the same location on my desk that he is gazing at on his. Our gaze traverses the distance, thereby strengthening the sense of presence. What is central to this example is the contrast between the simplicity and naturalness of the environment and the potency of its functionality. In keeping with the principle of invisibility, a powerful, non-intrusive work situation has been created.

Recapitulation: From the Macro to the Micro

Throughout this chapter, I have referred to the notion of “task space” and “person space” and emphasized the importance of keeping them separate. Before concluding, I want to drill down on this because something is missing, a bridge, which emerges if we look at things at finer granularity.

Fig. 13.8 Referencing with Shadows. In this example from Tang and Minneman's 1991 Videowhiteboard, one sees the remote person, but not like we have seen in the other examples. Here what is important is the relationship between shadow relative to the work.

The heart of what I am getting at is that there is a place where the space of the person and the task overlap, and this is what I am going to call *reference space.* This is perhaps best explained by referring to Tang and Minneman's 1991 *Videowhiteboard* system, illustrated in Fig. 13.10. Contrast this to the photograph in Fig. 13.8. In both cases one can see the remote participant(s) and the work being done - in both cases on a large rear projection screen. But here the similarities disappear rather quickly. In Fig. 13.10, one can see no details of the remote person's face, such as their eyes or where they are looking. On the other hand, in Fig. 13.8, the only way that people can point or gesture is with a single point, controlled by a mouse or stylus. This restricts them to the gestural vocabulary of a fruit fly. What a contrast to Fig. 13.10 where one has the full use or both hands and the body to reference aspects of the work through gestures. As well, the sharpness and contrast of the shadows provide strong cues that help one anticipate what the remote person is about to do, and where.

What Tang and Minneman gave us was a reminder that rich body language and gesture are important components of collaborative action. I would add is that it is useful to consider such overlaying of the human gesture on the work surface as a separate notion, distinct from – but complementary to – task space and person space. This is something that is starting to get traction in more recent work, such as C-Slate (Izadi, *et al* 2007), among others.

To summarize, I would identify three distinct types of spaces that need to be considered at the micro level of collaboration:

1. *Person space:* this is the space where one reads the cues about expression, trust, gaze. It is where the voice comes from, and where you look when speaking to someone.
2. *Task space:* this is the space where the work appears. If others can see it, it is shared. If not, it is private. Besides viewing, this is the space where one does things, such as marking or creating. One changes things here.
3. *Reference space:* the is the space within which the remote party can use body language to reference the work – things like pointing, gesturing. It is also the channel through which one can sense proximity, approach, departure, and anticipate intent. Like the task space, there are different types. In this case, the types vary according to richness, with a

telepointer being pretty low on the scale, and the shadows of Tang and Minneman fairly high.

The reason that I make these distinctions is to emphasize that person space and reference space need not be the same thing. The most interesting evidence of this is that one can have both, yet disconnected. For example, in Fig. 13.10, if there was a high quality video image of the remote person's head hanging like a rear-view mirror of the left of the frame, there would be no disorientation in seeing the shadow of the body, including the head superimposed on the work surface, and the video image of the face to the left. The reason is that they serve different purposes, even if they are the same person. In some ways, this is not dissimilar to our ability to "point" with a mouse, where our hand is on the desk, while our eye tracks the cursor.

The importance of this is directly related to the importance of being able to effectively frame and make reference to things that we are working on in the everyday world. It is a rich form of communication, and telepresence systems that do not support it well will be impoverished as a result – no matter how good the audio and video might otherwise be.

Summary and Conclusions

Even with the best design, the systems that we create impose a load on our users. The problem is, in all likelihood, the reason that they adopted the technology in the first place was because they were already pushing the limits of what they could handle. Hence, our job is to ensure that we take best advantage of the skills that they already have, so that we minimize things that they have to do or learn. In the world of telepresence and mediaspaces, the knowledge and skills that offer the lowest hanging fruit, are those associated with our collective understanding of place, location, distance, function, and meaning. And yet, the potential of these same skills is all too often ignored in our designs. The hope is that the work described in this chapter helps shed some light on this potential.

Acknowledgments The ideas developed in this essay have evolved over countless discussions with colleagues at Rank Xerox EuroPARC, Xerox PARC and the Ontario Telepresence Project. To all of those who have helped make these such stimulating environments, I am very grateful. I would like to especially acknowledge the contributions of Abi Sellen, Sara Bly, Steve Harrison, Mark Weiser, Brigitta Jordan and Bill Gaver. Finally, I would like to acknowledge the contribution of my wife, Elizabeth Russ, who made many helpful comments on the manuscript. The research discussed in this paper has been supported by the Ontario Telepresence Project, Xerox PARC and the Natural Sciences and Engineering Research Council of Canada. This support is gratefully acknowledged.

References

Bly, S., Harrison, S. & Irwin, S. (1993). Media Spaces: bringing people together in a video, audio and computing environment. *Communications of the ACM*, 36(1), 28-47.

Buxton, W. (2006). Space-Function Integration and Ubiquitous Media. In M. Shamiyeh (Ed.). *Towards and Interactive and Integrative Design Process*. Linz, Austria: DOM Publications, 248-271.

Buxton, W. (1992). Telepresence: Integrating Shared Task and Person Spaces . *Proceedings of Graphics Interface '92,* 123-129. Earlier version appears in Proceedings of Groupware '91, Amsterdam, Oct. 29, 1991, 27-36.

Buxton, W. (1995). Integrating the Periphery and Context: A New Model of Telematics. *Proceedings of Graphics Interface '95*, 239-246.

Buxton, W. (1997). Living in Augmented Reality: Ubiquitous Media and Reactive Environments. In K. Finn, A. Sellen & S. Wilber (Eds.). *Video Mediated Communication*. Hillsdale, N.J.: Erlbaum, 363-384. An earlier version of this chapter also appears in *Proceedings of Imagina '95*, 215-229.

Buxton, W. & Moran, T. (1990). EuroPARC's Integrated Interactive Intermedia Facility (iiif): Early Experience. In S. Gibbs & A.A. Verrijn-Stuart (Eds.). *Multi-user interfaces and applications, Proceedings of the IFIP WG 8.4 Conference on Multi-user Interfaces and Applications*, Heraklion, Crete. Amsterdam: Elsevier Science Publishers B.V. (North-Holland), 11-34.

Buxton, W., Sellen, A. & Sheasby, M. (1997). Interfaces for Multiparty Videoconferencing. In K. Finn, A. Sellen & S. Wilber (Eds.). *Video Mediated Communication.* Hillsdale, N.J.: Erlbaum, 385-400.

Cooperstock, J., Fels, S., Buxton, W. & Smith, K.C. (1997). Reactive environments: Throwing away your keyboard and mouse, *Communications of the ACM*, 40(9), 65-73.

Dourish, P. (2006). Re-Space-ing Place: "Place" and Space" Ten Years On. *Proceedings of Computer Supported Collaborative Work '06, 299 – 308.*

Fields, C.I. (1983). Virtual space teleconference system. *United States Patent 4,400,724*, August 23, 1983.

Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K. & Buxton, W. (1992). Realizing a video environment: EuroPARC's RAVE System. *Proceedings of CHI '92*, *ACM Conference on Human Factors in Software* 27-35.

Harrison, S. & Dourish, P. (1996). Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems. *Proceedings of Computer Supported Collaborative Work '96,* 67 – 76.

Ishii, H., Kobayashi, M. & Grudin, J. (1992). Integration of inter-personal space and shared workspace: Clearboard design and experiments. *Proceedings of CSCW '92*, 33 - 42.

Izadi,S., Agarwal, A., Criminisi, A., Winn, J. Blake, A. & Fitzgibbon,, A. (2007). C-Slate: Exploring Remote Collaboration on Horizontal Multi-touch Surfaces. *Proc. IEEE Tabletop 2007*, 3 – 10.

Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T. & Welleman, B. (1991). Experiences in the use of a media space. *Proceedings of CHI '91, ACM Conference on Human Factors in Software*, 203-208.

Sellen, A. (1992). Speech patterns in video mediated conferences. *Proceedings of CHI '92, ACM Conference on Human Factors in Software*, 49-59.

Sellen, A., Buxton, W. & Arnott, J. (1992). Using Spatial Cues to Improve Videoconferencing. *Proceedings of CHI '92, ACM Conference on Human Factors in Software*, 651-652. Also videotape in CHI '92 Video Proceedings.

Stults, R. (1986). Media Space. *Systems Concepts Lab Technical Report*. Palo Alto, CA: Xerox PARC.

Tang, J.C. & Minneman, S.L. (1991). Videowhiteboard: Video shadows to support remote collaboration. *Proceedings of the ACM-SIGCHI Conference on Human Factors in Computing Systems (CHI'91),* 315–322.

Weiser, M. (1991). The Computer for the 21st Century. *Scientific American*, 265(3), 94-104.