

A Wearable Spatial Conferencing Space

M. Billingham^α, J. Bowskill^β, M. Jessop^β, J. Morphet^β

^αHuman Interface Technology Laboratory
University of Washington
Box 352-142
Seattle, WA 98195, USA

^βAdvanced Perception Unit
BT Laboratories
Martlesham Heath
Ipswich, IP5 3RE, United Kingdom

grof@hitl.washington.edu {jerry.bowskill, jason.morphett, mark.jessop}@bt-sys.bt.co.uk

Abstract

Wearable computers provide constant access to computing and communications resources. In this paper we describe how the computing power of wearables can be used to provide spatialized 3D graphics and audio cues to aid communication. The result is a wearable augmented reality communication space with audio enabled avatars of the remote collaborators surrounding the user. The user can use natural head motions to attend to the remote collaborators, can communicate freely while being aware of other side conversations and can move through the communication space. In this way the conferencing space can support dozens of simultaneous users. Informal user studies suggest that wearable communication spaces may offer several advantages, both through the increase in the amount of information it is possible to access and the naturalness of the interface.

1: Introduction

One of the broad trends emerging in human-computer interaction is the increasing portability of computing and communication facilities. However it remains an open question as to how computing can best be used to aid mobile communication.

Wearable computers are the most recent generation of portable machines. Worn on the body, they provide constant access to computing and communications resources. In general, a wearable computer may be defined as a computer that is subsumed into the personal space of the user, controlled by the wearer and has both

operational and interactional constancy, i.e. is always on and always accessible [1]. Wearables are typically composed of a belt or back pack PC, see-through or see-around head mounted display (HMD), wireless communications hardware and an input device such as touchpad or chording keyboard. This configuration has been demonstrated in a number of real world applications including aircraft maintenance [2], navigational assistance [3] and vehicle mechanics [4]. In such applications wearables have dramatically improved user performance, halving task time in the case of vehicle inspection [4].

Many of the target application areas are those where the user could benefit from expert assistance. Network enabled wearable computers can be used as a communications device to enable remote experts to collaborate with the wearable user. In such situations the presence of remote experts have been found to significantly improve task performance [5], [6]. For example, Kuzuoka finds that a head mounted camera and display increases interaction efficiency in a remote collaborative object manipulation task [7]. However, most current collaborative wearable applications have only involved connections between one local and one remote user. The problem we are interested in is how a wearable computer can be used to support collaboration between multiple remote people. In particular we want to explore the following issues:

- What visual and audio enhancements can be used to aid communication?
- How can a collaborative communications space be created between users?
- How can remote users be represented in a wearable computing environment?

These issues are becoming increasingly important as telephones incorporate more computing power and portable computers become more like telephones. A key issue is whether we need computer-mediated communication at all when a conference phone call may be just as effective. Prior work in the fields of teleconferencing and computer supported collaborative work addresses this question.

2: Background

Research on the roles of audio and visual cues in teleconferencing has produced mixed results. There have been many experiments conducted comparing face-to-face, audio and video, and audio only communication conditions. Sellen summarizes these by reporting that the main effect on collaborative performance is due to whether the collaboration was technologically mediated or not, not on the type of technology mediation used [8]. While people generally do not prefer the audio only condition, they are often able to perform tasks as effectively as in the audio and video condition, although in both cases they perform worse than face-to-face collaboration. Naturally this varies somewhat according to task. While face-to-face interaction is no better than speech-only for cognitive problem solving tasks [9], visual cues can be important in tasks requiring negotiation [10]. In general, the usefulness of video for transmitting non-verbal cues may be overestimated and video may be better used to show the communication availability of others or views of shared workspaces [11]. Even when users attempt non-verbal communication in a video conferencing environment their gestures must be exaggerated to be recognized as the equivalent face-to-face gestures [12].

Based on these results, and the fact that speech is the critical medium in teleconferencing experiments [13], it may be thought that audio alone should be suitable for a creating a shared communication space. An example of this, Thunderwire [14], was a purely audio system which allowed high quality audio conferencing between multiple participants at the flip of a switch. In a 3 month trial Hindus et. al. found that audio can be sufficient for a usable communication space. However several major problems were observed:

- Users were not able to easily tell who else was within the space.
- Users were not able to use visual cues to determine other's willingness to interact.

With more users it becomes increasingly difficult to discriminate between speakers and there is a higher

incidence of speaker overlap and interruptions. These problems are typical of audio only spaces and suggest that while audio may be useful for small group interactions, it becomes less usable the more people present.

These shortcomings can be overcome through the use of visual and spatial cues. In face-to-face conversation, speech, gesture, body language and other non-verbal cues combine to show attention and interest. Simple versions of these cues can be replicated in desktop video conferencing. For example, the Passepartout [15] enhanced desktop conferencing tool includes visual representation of a conference table alongside shared documents and a text chat facility. The conference table includes icons of those people within the conference and simple cues, such as microphone on/off, which allow a participant's activity or interest to be inferred. However the absence of spatial cues in most video conferencing systems means that users often find it difficult to know when people are paying attention to them, to hold side conversations, and to establish eye contact [16].

Several video conferencing systems have attempted to provide spatial cues. The Hydra system uses multiple small monitors, one for each participant, positioned about the local user [17]. The user can easily attend to individual participants by turning to face the appropriate monitor and side conversations can be supported. The MAJIC system uses several wall projectors and a one way transmissive screen to create the illusion of several remote life-sized participants seated around the same real table [18]. Users can make eye contact and conduct parallel conversations. The MPEC prototype also supports spatial video conferencing and multi-party eye contact [19]. However a common disadvantage of these systems is that the users cannot control remote camera position and so their viewpoint and spatial relationships to the other participants is fixed, unlike in face-to-face collaboration. There are also many technical problems to be overcome before such systems scale to support large groups of participants.

Virtual reality can provide an alternative medium that allows groups of people to share the same communications space. British Telecom has demonstrated virtual conferencing in which many users can be represented as life like virtual avatars of themselves within virtual rooms [20]. Users can freely move through the space setting their own viewpoints and spatial relationships. In collaborative virtual environments (CVEs) spatial visual and audio cues can combine in natural ways to aid communication [21]. The well known "cocktail-party" effect shows that people can easily

monitor several spatialized audio streams at once, selectively focusing on those of interest [22], [23]. Even a simple virtual avatar representation and spatial audio model enables users to discriminate between multiple speakers [24]. Spatialized interactions are particularly valuable for governing interactions between large groups of people; enabling crowds of people to inhabit the same virtual environment and interact in a way impossible in traditional video or audio conferencing [25].

3: A Wearable Communication Space

The results in the previous section suggest that an ideal wearable communications space should have three elements:

- High quality audio communication
- Visual representations of the collaborators
- An underlying spatial metaphor

One of the most important aspects of creating a collaborative communication interface is the visual and audio presentation of information. Most current wearable computers use see-through or see-around monoscopic head mounted displays with stereo headphones. With these displays information can be presented in a combination of three ways:

Head-stabilized - information is fixed to the users viewpoint and doesn't change as the user changes viewpoint orientation or position.

Body-stabilized - information is fixed relative to the users body position and varies as the user changes viewpoint orientation, but not as they change position. This requires the users viewpoint orientation to be tracked.

World-stabilized - information is fixed to real world locations and varies as the user changes viewpoint orientation and position. This requires the users viewpoint position and orientation to be tracked.

Body and World stabilized information display is attractive for a number of reasons. As Reichlen [26] demonstrates, a body-stabilized information space can overcome the resolution limitations of head mounted displays. In his work a user wears a head mounted display while seated on a rotating chair. By tracking head orientation the user experiences a hemispherical information surround - in effect a "hundred million pixel display". World-stabilized information presentation enables annotation of the real world with context dependent visual and audio data, creating information enriched environments [27]. This increases the

intuitiveness of real world tasks. Despite these advantages, most current wearables only use head-stabilized information display.

In our work we have chosen to begin with the simplest form of body-stabilized display; one which uses one degree of orientation to give the user the impression they are surrounded by a virtual cylinder of visual and auditory information. Figures 1.0a and 1.0b contrast this with the traditional head stabilized wearable interface.




Figure 1.0a Head Stabilized Information Display




Figure 1.0b One Degree of Freedom Body-Stabilized Display

When using a head mounted display to navigate a cylindrical body-stabilized space, only the portion of the information space in its field of view can be seen. There are two ways the rest of the space can be viewed; by rotating the information space about the user's head, or tracking the user's head orientation as they look around the space. The first requires no additional hardware and

can be done by mapping mouse, switch or voice input to direction and angle of rotation, while the second requires only a simple one degree of freedom tracker. The minimal hardware requirements make cylindrical spatial information displays particularly attractive. The cylindrical display is also very natural to use since most head and body motion is about the vertical axis, making it very difficult for the user to become disoriented.

In a previous paper we have found that users can locate information more rapidly with this type of information display than the more traditional head-stabilized wearable information space [28]. Sawhney and Schmandt also demonstrate how body-stabilized spatial audio can improve access to audio information on a wearable platform, allowing a user to browse up to three simultaneous audio streams [29].

With this display configuration a wearable conferencing space could be created that allows remote collaborators appear as virtual avatars distributed about the user (figure 2.0). The avatars could be live video streams and as they speak their audio streams spatialized in real time so that they appear to emit from the corresponding avatar.




Figure 2.0 A Spatial Conferencing Space.

Just as in face-to-face collaboration, users could turn to face the collaborators they wanted to talk to while still being aware of the other conversations taking place. The user could also move about the space enabling them to choose their own viewpoint and the spatial relationships between the collaborators. In this way the space could support dozens of simultaneous users, similar to current collaborative virtual environments. Since the displays are see-through or see-around the user could also see the real

world at the same time, enabling the remote collaborators to help them with real world tasks. These remote users may also be using wearable computers and head mounted displays or could be interacting through a desktop workstation. The wearable conferencing space would also allow the faces of remote users to appear life-size, a crucial factor for establishing equal relationships in remote collaboration [30].

The technical requirements for such a conferencing space place it several years in the future, however in the remainder of this paper we describe a prototype we have developed which has many of the same features.

4: Implementation

Our research is initially focused on collaboration between a single wearable computer user and several desktop PC users. This situation might be encountered when a wearable user in the field is requesting help from remote deskbound experts. The aim is to develop a wearable interface to support medium sized meetings (5-6 people) in a manner that is natural and intuitive to use.

4.1: Hardware

The wearable computer we use is a custom built 586 PC 104 based computer with 20mb of RAM running Windows 95. Figure 3.0 shows a user wearing the display and computer.



Figure 3.0 The Wearable Hardware.

A hand held Logitech wireless radio trackball with three buttons is used as the primary input device. The display is a pair of Virtual i-O iglasses! converted into a monoscopic display by the removal of the left eyepiece.

The Virtual i-O head mounted display can either be used in see-through or occluded mode, has a resolution of 262 by 230 pixels and a 26-degree field of view. The iglasses! also have a sourceless two-axis inclinometer and a magnetometer used as a three degree of freedom orientation tracker. A BreezeCom wireless LAN is used to give 2mb/s Internet access up to 500 feet from a base station. The wearable also has a soundBlaster compatible sound board with headmounted microphone. The desktop PCs are standard Pentium class machines with Internet connectivity and sound capability.

4.2: The Wearable Interface

Our wearable computer has no graphics acceleration hardware and limited wireless bandwidth so the interface is deliberately kept simple. The conferencing space runs as a full screen application that is initially blank until remote users connect. When users join the conferencing space they are represented by blocks with 128x128 pixel texture mapped static pictures of themselves on them. Each user determines the position and orientation of their own avatar in space, which changes as they move or look about environment. Although the resolution of the images is crude it is sufficient to identify who the speakers are and more importantly their spatial relationship to the wearable user. It is hoped that in the near future wearable computer CPU power and wireless bandwidth will be sufficient to support real time video texture mapping.

The wearable user has their head tracked so they can simply turn to face the speakers they are interested in. Users can also navigate through the space; by rolling the trackball forwards or backwards their viewpoint is moved forwards or backwards along the direction they are looking. Since the virtual images are superimposed on the real world, when the user rolls the trackball it appears to them as though they are moving the virtual space around them, rather than navigating through the space. Users are constrained to change viewpoint on the horizontal plane, just as in face-to-face conversations. The two different navigation methods (trackball motion, head tracking), match the different types of motion used in face to face communication; walking to join a group for conversation, and body orientation changes within a conversational group.

A radar display shows the location of the other users in the conferencing space, enabling users to find each other easily. Figure 4.0 shows the wearable interface from the wearable user's perspective. The interface was developed using Microsoft's Direct3D, DirectDraw and DirectInput libraries from the DirectX suite.

The wearable interface also supports 3D spatialized Internet telephony. When users connect to the conferencing space their audio is broadcast to all the other users in the space. This is spatialized according to the distance and direction between speaker and listener. As users face or move closer to different speakers the speaker volume changes due to the sound spatialisation. Since the speakers are constrained to remain in the same plane as the listener the audio spatialisation is considerably simplified. Audio culling is also used so that only the audio streams from the speakers closest to the listener are broadcast and spatialized. This significantly reduces the CPU load. The conferencing space uses custom developed telephony libraries that incorporate the Microsoft Direct Sound libraries.




Figure 4.0 The User's View of the Wearable Conferencing Space

4.3: The Desktop interface

Users at a desktop workstation interact with the conferencing space through a similar interface as the wearable user, although in this case the application runs as a Windows application on the desktop. Users navigate through the space using the mouse. Mouse movements rotate head position when the left mouse button is held down, otherwise they translate the user backwards and forwards in space. Mapping avatar orientation to mouse movement means that the desktop interface is not quite as intuitive as the wearable interface. Users at the desktop machine wear headmounted microphones to talk into the conferencing space and listen through stereo headphones. Just as with the wearable interface, desktop users are aware of the spatial relationships between participants. When a participant turns and talks to someone else, the desktop user sees their avatar turn and face the person they're talking to.

5: Distributed Software Architecture

The wearable and desktop interfaces are based around custom libraries for collaborative virtual environments being developed at British Telecom. When the wearable and desktop client applications are run TCP/IP multicast groups are created that enable the clients to communicate with each other through multicast sockets. The multicast protocol is an efficient mechanism for broadcasting data to multiple network nodes [31] and has been shown to scale well in large CVEs [32].

Communications within the conferencing space are routed through one of two multicast groups, as shown in figure 5.0. One is for transformational data representing an avatar's position and orientation plus any messaging data and the second for audio data. When users connect to the communication space they are assigned a unique identification tag, (ID). As a user moves through the space, their avatar's ID tag, positional and orientation information is broadcast onto the transformation multicast group. This transformational data flows at a rate of 10.0Kb/s per user. When received by each client in the group it is used to update the relevant avatar's position and orientation. The transformational information is also used in spatialising the user's audio stream relevant to the receiving user's position.

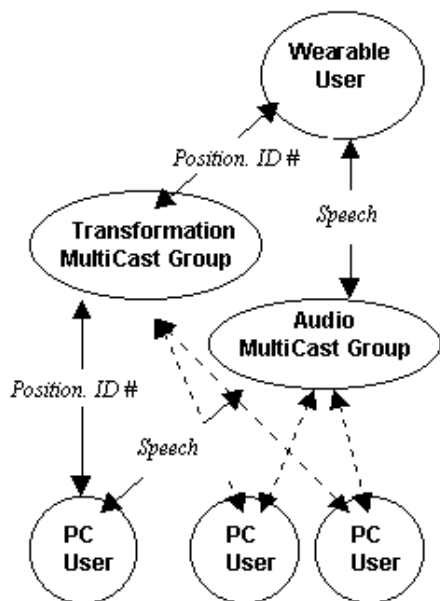


Figure 5.0 Distributed Software Architecture

Similarly when a user speaks, their speech is digitized and broadcast to the audio multicast group. When received by the other clients the senders IP address

identifies the avatar that the audio belongs to and its position and orientation. The audio is then spatialized in real time. The audio is implemented utilizing Microsoft's Direct Sound technology. This allows for the capture of audio to a buffer, which can then be broadcast over the audio multicast group. Once received at a client computer the buffer can be played back through Direct Sound. In order for the audio to operate in full duplex mode it has to be captured at a rate of 8 bit 22KHz, resulting in a data rate 172Kb/s. All connections to the multicast groups are bi-directional and users can connect and disconnect at will without affecting other users in the conferencing space.

6: Initial User Experiences

In developing a wearable conferencing space we set out to explore the usefulness of spatial visual and audio cues compared to traditional portable communications devices, namely audio only collaboration with a mobile phone. We are in the process of conducting user trials to evaluate how the use of spatialized audio and visual representations affects communication between collaborators. Preliminary informal trials have found the following results:

Users are able to easily discriminate between three simultaneous speakers when their audio streams are spatialized, but not when non-spatialized audio is used. It is expected that this effect will become even more noticeable as the number of simultaneous participants is increased.

Participants preferred seeing a visual representation of their collaborators as opposed to just hearing their speech. Even though it was relatively poor quality the visual representation enabled them to see who is connected and the spatial relationship of the speakers. This allowed them to use some of the non-verbal cues commonly used in face-to-face communication such as gaze modulation and body motion. The radar display was useful for finding collaborators that were far away and barely visible.

Users found that they could continue doing real world tasks while talking to collaborators in the conferencing space and it was possible to move the conferencing space with the trackball so that collaborators weren't blocking critical portions of the users field of view.

The interface is easy and intuitive to use, although using the head tracking on the wearable was easier than the mouse only desktop interface.

However, as more users connect to the conferencing space the need to spatialize multiple audio streams puts a

severe load on the CPU, slowing down the graphics and head tracking. This makes it difficult for the wearable user to conference with more than two or three people simultaneously. This problem will be reduced as faster CPUs and hardware support for 3D graphics become available for wearable computers. More severe spatial culling of the audio streams could also be used to overcome this limitation, through the coagulation of selected streams into a single spatial location, or removing the audio altogether.

7: Conclusions

We have presented a prototype wearable communication space that uses spatial visual and audio cues to enhance communication between remote groups of people. Our interface shows what is possible when computing and communications facilities are coupled together on a wearable platform. Preliminary results have found that users prefer using both the audio and visual cues together and that spatialized audio makes it easy for users to discriminate between speakers. This suggests that for some applications wearable computers may provide a useful alternative to traditional audio-only communication devices. We are currently conducting formal user studies to confirm these results and evaluate the effect of spatial cues on communication patterns. In the future we plan to investigate how the presence of spatialized video can further enhance communication. We will incorporate live video texture mapping into our interface, enabling users to see their remote collaborators as they speak. This will also allow users to send views of their workspace, improving collaboration on real-world tasks.

We believe that a wearable communications space can be used to support numerous collaborative applications in which some participants are either not sitting at desks or need mobility. A shared virtual environment facilitates audiovisual communications for groups of people with the added potential for embedded graphical, textual or audio information. A specific trait that we believe to be particularly important is that in an augmented communications space it is possible for the user to form effective cognitive maps by associating the annotated information with physical objects within their surroundings. We have demonstrated a body-stabilized system in which the communications space is located relative to the user themselves. With the ability to explicitly position object in the communications space relative to the user's real world location an exciting range of applications become possible. A user could, for example, choose to view avatars of conference participants overlaid and attached to a physical notice

board. This represents a powerful vision of conferencing for all platforms; with remote video conferencing participants not in separate windows on a screen but spread around the user's environment, positioned in space where the user prefers.

8: Acknowledgements

We would like to thank our colleagues at British Telecom and the HIT Lab for many insightful and productive conversations, Nick Dyer for producing the renderings used in some of the figures and the anonymous reviewers for their useful comments.

9: References

- [1] Mann, S. Smart Clothing: The Wearable Computer and WearCam. *Personal Technologies*, Vol. 1, No. 1, March 1997, Springer-Verlag.
- [2] Esposito, C. Wearable Computers: Field-Test Results and System Design Guidelines. In *Proceedings Interact '97*, July 14th-18th, Sydney Australia.
- [3] Feiner, S., MacIntyre, B. Hollerer, T. A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment. In *Proceedings of the International Symposium on Wearable Computers*, Cambridge, MA, October 13-14, 1997, Los Alamitos: IEEE Press, pp. 74-81.
- [4] Bass, L., Kasabach, C., Martin, R., Siewiorek, D., Smailagic, A., Stivoric, J. The Design of a Wearable Computer. In *Proceedings of CHI '97*, Atlanta, Georgia. March 1997, New York: ACM, pp.139-146.
- [5] Siegal, J., Kraut, R., John, B., Carley, K. An Empirical Study of Collaborative Wearable Computer Systems,. In *Proceedings of CHI 95 Conference Companion*, May 7-11, Denver Colorado, 1995, ACM: New York, pp. 312-313.
- [6] Kraut, R., Miller, M., Siegal, J. Collaboration in Performance of Physical Tasks: Effects on Outcomes and Communication. In *Proceedings of CSCW '96*, Nov. 16th-20th, Cambridge MA, 1996, New York, NY: ACM Press.
- [7] Kuzuoka, H. Spatial Workspace Collaboration: A Shared View Video Support System for Remote Collaboration. In *Proceedings of CHI '92 Human Factors in Computing Systems*, Monterey, CA, May 3-7, 1992, ACM:New York, pp. 533-540.
- [8] Sellen, A. Remote Conversations: The effects of mediating talk with technology. *Human Computer Interaction*, 1995, Vol. 10, No. 4, pp. 401-444.
- [9] Williams, E. Experimental Comparisons of Face-to-Face and Mediated Communication. *Psychological Bulletin*, 1997, Vol 16, pp. 963-976.

- [10] Chapanis, A. Interactive Human Communication. *Scientific American*, 1975, Vol. 232, pp36-42.
- [11] Whittaker, S. *Rethinking Video as a Technology for Interpersonal Communications: Theory and Design Implications*. Academic Press Limited, 1995.
- [12] Heath, C., Luff, P. Disembodied Conduct: Communication Through Video in a Multimedia Environment. In *Proceedings of CHI '91 Human Factors in Computing Systems*, 1991, New York, NY: ACM Press, pp. 99-103.
- [13] Whittaker, S., O'Connell, B. The Role of Vision in Face-to-Face and Mediated Communication. In *Video-Mediated Communication*, Eds. Finn, K., Sellen, A., Wilbur, S. Lawrence Erlbaum Associates, New Jersey, 1997, pp. 23-49.
- [14] Hindus, D., Ackerman, M., Mainwaring, S., Starr, B. Thunderwire: A Field study of an Audio-Only Media Space. In *Proceedings of CSCW '96*, Nov. 16th -20th, Cambridge MA, 1996, New York, NY: ACM Press.
- [15] Russ, M. Desktop Conversations - The Future of Multimedia Conferencing. *BT Technology Journal*, Vol. 14, No. 4, October 1997, pp. 42-50.
- [16] Sellen, A. Speech Patterns in Video-Mediated Conversations. In *Proceedings CHI '92*, May 3-7, 1992, ACM: New York, pp. 49-59.
- [17] Sellen, A., Buxton, B. Using Spatial Cues to Improve Videoconferencing. In *Proceedings CHI '92*, May 3-7, 1992, ACM: New York, pp. 651-652.
- [18] Okada, K., Maeda, F., Ichikawa, Y., Matsushita, Y. Multiparty Videoconferencing at Virtual Social Distance: MAJIC Design. In *Proceedings of CSCW '94*, October 1994, New York: ACM, pp. 385-393.
- [19] DeSilve, L., Tahara, M., Aizawa, K., Hatori, M. A Teleconferencing System Capable of Multiple Person Eye Contact (MPEC) Using Half Mirrors and Cameras Placed at Common Points of Extended Lines of Gaze. *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 5, No. 4, August 1995, pp.268-277.
- [20] Mortlock, A., Machin, D., McConnell, S., Sheppard, P. Virtual Conferencing. *BT Technology Journal*, Vol. 14, No. 4, October 1997, pp. 120-129.
- [21] Benford, S. and Fahlen, L. A Spatial Model of Interaction in Virtual Environments. In *Proceedings of Third European Conference on Computer Supported Cooperative Work (ECSCW'93)*, Milano, Italy, September 1993.
- [22] Bregman, A. *Auditory Scene Analysis: The Perceptual Organization of Sound*. MIT Press, 1990.
- [23] Schmandt, C., Mullins, A. AudioStreamer: Exploiting Simultaneity for Listening. In *Proceedings of CHI '95 Conference Companion*, May 7-11, Denver Colorado, 1995, ACM: New York pp. 218-219.
- [24] Nakanishi, H., Yoshida, C., Nishimura, T., Ishida, T. FreeWalk: Supporting Casual Meetings in a Network. In *Proceedings of CSCW '96*, Nov. 16th -20th, Cambridge MA, 1996, New York, NY: ACM Press, pp. 308-314.
- [25] Benford, S., Greenhalgh, C., Lloyd, D. Crowded Collaborative Virtual Environments. In *Proceedings of CHI '97*, Atlanta, Georgia. March 1997, New York: ACM, pp.59-66
- [26] Reichlen, B. SparcChair: One Hundred Million Pixel Display. In *Proceedings IEEE VRAIS '93*. Seattle WA, September 18-22, 1993, IEEE Press: Los Alamitos, pp. 300-307.
- [27] Rekimoto, J., Nagao, K. The World through the Computer: Computer Augmented Interaction with Real World Environments. In *Proceedings of User Interface Software and Technology '95 (UITS '95)*, November 1995, New York: ACM, pp. 29-36.
- [28] Billingham, M., Bowskill, J., Dyer, N., Morphett, J. An Evaluation of Wearable Information Spaces. In *Proceedings of IEEE VRAIS '98*, Atlanta, Georgia, March 14th-18th, 1998, IEEE Computer Society Press, Los Alamitos, CA.
- [29] Sawheny, N., Schmandt, C. Design of Spatialized Audio in Nomadic Environments. In *Proceedings of the International Conference on Auditory Display (ICAD '97)*, Palo-Alto, November 5th, 1997.
- [30] King, J. Human Computer Dyads? A Survey of Nonverbal Behavior in Human-Computer Systems. In *Proceedings of the Workshop on Perceptual User Interfaces (PUI '97)*, Banff, Canada, Oct. 19-21, IEEE Computer Society Press, Los Alamitos, CA, pp.54-55.
- [31] Kumar, V. *Mbone: Interactive Multimedia on the Internet*. New Riders, Indianapolis, Indiana, 1995.
- [32] Macedonia, M., Zyda, M., Pratt, D. Exploiting Reality with Multicast Groups: A Network Architecture for Large-Scale Virtual Environments. In *Proceedings of the IEEE VRAIS '95 Conference*. IEEE Computer Society Press, Los Alamitos, CA, March 1995, pp. 2-10.