Crossing the Chasm

Mark Billinghurst

Human Interface Technology Laboratory,
University of Washington
Box 352-142, Seattle, WA 98195 JAPAN
grof@hitl.washington.edu

Abstract

Traditional computer interfaces often create a separation between physical and digital spaces. This prevents people from working as easily in virtual worlds as they do in the real world, or from moving seamlessly between reality and virtual reality. In this paper we describe how this chasm can be crossed though the use of an interface metaphor, Tangible Augmented Reality, that combines elements from both the physical and digital domains.

Key words: Augmented Reality, Virtual Reality, CSCW, Tangible User Interfaces.

1. Introduction

As computers become more and more powerful while at the same time decreasing in size, an important question is not what we can use the processing power for, but how we can best interact with the machine. In the near future, through a combination of high-bandwidth wireless networks, cheap powerful processors and pervasive computing technologies we will be able to have all the CPU power that we may need for almost any task. Like the mythical "Ether" of a century ago computing power may seen to pervade the very atmosphere itself.

Ironically, the trend toward pervasive computing increases availability to processing power while at the same time reduces access through outmoded interface design. Traditional interface metaphors do not scale across non-traditional displays or input devices and may be flawed in the assumptions they make about how the computer is going to be used. For example, the dominant WIMP (Windows, Icons, Menus, Pointers) metaphor is based on the inherent assumption that the users main task is interacting with the computer, however in many mobile applications interaction with the real world is the primary goal.

One of the greatest limitations of traditional interface design is the separation between the real and digital worlds. A person can interact with a computer thorough the familiar desktop interface, but apart from keyboard and mouse input the computer remains oblivious to events in the real world. Similarly, digital information can be manipulated without ever leaving a mark in the physical world. The computer cannot disappear into the real world until it becomes aware of it, so there is a chasm that exists between physical and digital reality that must be bridged before pervasive computing can be truly useful.

Recognizing the need to blur the line between atoms and bits, researchers have developed interfaces in which the computer vanishes into familiar real world objects or the environment. We have seen the development of the Digital Desktop [16], wearable computers [14], ubiquitous computing [15] and tangible computing [9] among others. The paradigm shift represented by these efforts is so significant that Weiser labeled ubiquitous computing as "The Computer for the 21st Century".

One of the more interesting ways to blend the real and digital domains is through the use of Augmented Reality (AR)). This is where three-dimensional computer graphics are superimposed over real objects, and are typically seen through head-mounted or handheld displays. However, the AR field has been primarily concerned with "..considering purely visual augmentations" [9] and while great advances have been made in AR display technologies and tracking techniques, *interaction* with AR environments has been usually limited to either passive viewing or simple browsing of virtual information registered to the real world. Few systems provide tools that let the user interact, request or modify this information effectively and in real time.

In this paper we show how the chasm between the real and digital worlds can be spanned by using an interface metaphor we call Tangible Augmented Reality (Tangible AR). This is an interface design methodology in which combines elements of Augmented Reality display with tangible object interaction. Tangible AR interfaces are those in which 1) each virtual object is registered to a physical object and 2) the user interacts with virtual objects by manipulating the corresponding tangible objects. Thus the physical objects and interactions are equally as important as the virtual imagery.

In the remainder of this paper we explain the Tangible AR concept in more detail and then give examples of Tangible AR interfaces that others and we have developed.

2. Tangible User Interfaces

As mentioned above, several research groups have begun to explore Tangible User Interface (TUI) in which real world objects are used as computer input and output devices, or as Hirishi Ishii puts it "by coupling digital information to everyday physical objects and environments" [9]. Tangible interfaces are extremely intuitive to use because physical object manipulations are mapped one-to-one to virtual object operations, and they follow a space-multiplexed input design [4].

In general input devices can be classified as either space- or time-multiplexed. With a space-multiplexed interface each function has a single physical device occupying its own space. In contrast, in a time-multiplexed design a single device controls different functions as different points in time. The mouse in a WIMP interface is a good example of a time-multiplexed device. Space-multiplexed devices are faster to use than time-multiplexed devices because users do not have to make the extra step of mapping the physical device input to one of several logical functions. In most manual tasks space-multiplexed devices are used to interact with the surrounding physical environment.

Although intuitive to use, with TUI interfaces information display can be a challenge. It is difficult to dynamically change an object's physical properties, so most information display is confined to image projection on objects or augmented surfaces. In those TUI that use three-dimensional graphics there is also often a disconnection between the task space and display space. For example, in the Triangles work [6], physical triangles are assembled to tell stories, but the visual representations of the stories are shown on a separate monitor distinct from the physical interface. Presentation and manipulation of 3D virtual objects on projection surfaces is difficult [5], particularly when trying to support multiple users each with independent viewpoints. Most importantly, because the information display is limited to a projection surface, users are not able to pick virtual images off the surface and manipulate them in 3D space as they would a real object.

So we see that current Tangible interfaces provide very intuitive manipulation of digital data, but limited support for viewing 3D virtual objects. In contrast, current AR interfaces provide an excellent interface for viewing virtual models, but limited support for interaction and space-multiplexed input devices. These two interface metaphors are complimentary; tangible interfaces offer seamless interaction but results in spatial discontinuities, while AR interfaces support spatially seamless workspaces but introduce discontinuities in interaction (table 1).

	AR Interfaces	Tangible Interfaces
Spatial Gap	None – interaction is everywhere	Yes – interaction is only on 2D surfaces
Interaction Gap	Yes – separate devices for physical and virtual objects	No – same devices for physical and virtual objects

Table 1: Complimentary nature of Augmented Reality and Tangible Interfaces

We believe that a promising new AR interface metaphor can arise from combining the enhanced display possibilities of Augmented Reality with the intuitive manipulation of Tangible User Interfaces. We call this combination Tangible Augmented Reality. In the next section we show how Tangible AR supports seamless interaction, and provide some design guidelines.

3. Tangible Augmented Reality

The goal of computer interfaces is to facilitate seamless interaction between a user and their computer-supported task. In this context, Ishii defines a seam as a discontinuity or constraint in interaction that forces the user to shift among a variety of spaces or modes of operation [8]. Seams that force a user to move between interaction spaces are called *functional* seams, while those that force the user to learn new modes of operation are *cognitive* seams.

In the previous section we described how Tangible User Interfaces provide seamless interaction with objects, but may introduce a discontinuity or *functional* seam between the interaction space and display space. In contrast most AR interfaces overlay graphics on the real world interaction space and so provide a spatially seamless display. However they often force the user to learn different techniques for manipulating virtual content than from normal physical object manipulation or use a different set of tools for interacting with real and virtual objects. So AR interfaces may introduce a *cognitive* seam.

A Tangible AR interface provides true spatial registration and presentation of 3D virtual objects anywhere in the physical environment, while at the same time allowing users to interact with this virtual content using the same techniques as they would with a real physical object. So an ideal Tangible AR interface facilitates seamless display and interaction, removing the functional and cognitive seams found in traditional AR and Tangible User Interfaces. This is achieved by using the design principles learned from TUI interfaces, including:

- The use of physical controllers for manipulating virtual content.
- Support for spatial 3D interaction techniques (such as using object proximity).
- Support for both time-multiplexed and space-multiplexed interaction.
- Support for multi-handed interaction.
- Support for matching the physical constraints of the object to the requirements of the interaction task.
- The ability to support parallel activity where multiple objects are being manipulated.
- Collaboration between multiple participants

AR interfaces that follow these design principles will provide completely seamless interaction with virtual content and so will be extremely intuitive to use. In the next section we describe some prototype Tangible AR interfaces.

3. Tangible AR Example Interfaces

In order to explore the Tangible AR design space we have developed or assisted in the development of the following prototype interfaces:

Space-Multiplexed Interfaces	**Time-Multiplexed Interfaces**
Shared Space: A collaborative game [1]	*VOMAR:* A scene assembly application [11]
Tiles: A virtual prototyping application [12]	*ARgroove:* A music performance interface [13]
The MagicBook: A transitional interfaces [2]	
AR PRISM: A geospatial visualization interface [7]	

In this section we briefly describe three of these interfaces, Tiles, VOMAR and the MagicBook, showing how the Tangible AR design principles have been applied. With all of these interfaces we use the ARToolKit computer vision based tracking library [10]. This software provides camera pose information from black square markers and tracking patterns.

3.1 Tiles – Space Multiplexed Interaction

Tiles is an AR authoring interface that explores how more complicated behaviors can be supported, including copying, pasting, deleting, and browsing virtual information in AR settings. In Tiles we explore space-multiplexed control by assigning different behaviors to different objects, creating tangible 3D widgets. We distribute functionality across tangible AR widgets (that we called *tiles*) letting the user to choose operation simply by picking a needed tile. The application domain is rapid prototyping for aircraft instrument panels. The interface consists of a metal whiteboard, a book, and two stacks of magnetic tiles (approximately 15cm x 15cm). Sitting in front of the whiteboard the user wears a lightweight high resolution Sony Glasstron HMD with a video camera attached (figure 1).

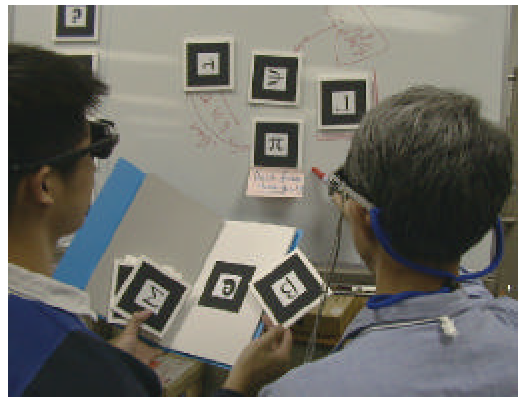

Figure 1: Using the Tiles Interface

The various tangible elements of the interface serve a different purpose. The whiteboard is the working space where users can layout virtual aircraft instruments. The book serves as a *menu object*, and when the user looks through its pages they will see a different virtual instrument model on each page. One stack of tiles serve as *data tiles* and shows no virtual content until virtual objects are copied onto them. The remaining tiles are *operator tiles* and are used to perform basic operations on the data tiles. There is a unique tile for each operation and currently supported operations include *deletion*, *copying* and a *help* function. Each of the operations tiles has a different three-dimensional virtual icon on them to show what their function is and tell them apart from the data tiles (figure 2).

Trashcan delete widget　　Talking head help widget

Figure 2: Virtual Widgets on Operator Tiles

Virtual images appear attached to the physical objects and can be picked up and looked at from any viewpoint. Interaction between objects is also based on physical proximity, however the operation that is invoked by bringing objects next to each other depends on their semantic. For example, to copy a virtual instrument from the menu book to an empty data tile, the tile is just placed by the appropriate book page. However, touching a data tile that contains a virtual instrument with the trashcan delete tile, removes the virtual instrument, while putting the help tile beside it displays a help message (2). Once virtual instruments have been placed on the data tiles, these can be attached to the whiteboard to layout a prototype virtual instrument panel (figure 3).

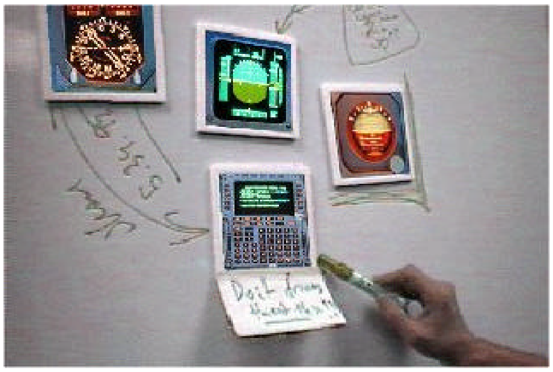

Figure 3: Virtual Instrument Panel view by the user

The two main features of this interface are the use of different shaped physical objects for different interface properties, and assigning different semantics to different objects. Supporting one interface function per object is similar to the interface models of desktop GUI interfaces, where each icon and tool has unique functionality. Despite this added functionality, the physical interactions are still based on object manipulation and proximity, showing that quite complex AR interfaces can be built from simple physical interactions. The use of different objects for different functions further emphasizes the space-multiplexed nature of the interface.

3.2 VOMAR – Time Multiplexed Interaction

The VOMAR (Virtual Object Manipulation in Augmented Reality) project explored how a time-multiplexed Tangible AR interface could be designed. It was primarily developed by Hirokazu Kato and is described in detail in another publication [11]. In this section we briefly review the VOMAR interface to show an example of a time-multiplexed interface with relatively complex physically based input.

VOMAR uses a single input device that allows the user to perform multiple different tasks in a virtual scene assembly application. To achieve this we explored how complex physical gestures can be used to support natural and effective interaction.

The physical components of the interface comprise a real book, a cardboard paddle the user holds in their hand, a large piece of paper and a lightweight HMD the user wears (figure 4a). The form of each of these objects reflects their function; the book serves as a container holding all the virtual models, the paddle is the main interaction device, and the large piece of paper the workspace.

Figure 4a The VOMAR interface

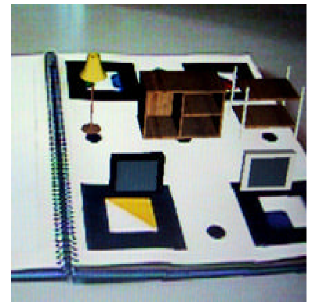

Figure 4b: Virtual Furniture

The application is layout of virtual furniture in a room, although the same interface can be applied to many domains. When the user opens the book on each of its pages they see a different set of virtual furniture, such as a set of chairs, rugs etc (fig 4b). The 3D virtual models appear exactly superimposed over the real book pages. Looking at the large piece of paper they see an empty virtual room. They can copy and transfer objects from the book to the room using the paddle (figure 5a, 5b).

Figure 5a Picking Furniture

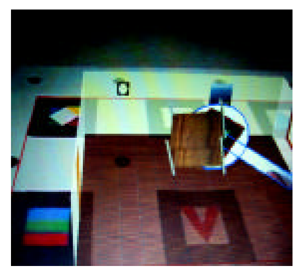

Figure 5b Placing Furniture

The paddle is the main interaction device and it is a simple object with an attached tracking symbol. It is designed to be used by either hand and allows the user to make static and dynamic gestures to interact with the virtual objects:

Static	Dynamic
1. Paddle proximity to object 2. Paddle tilt/inclination	1. Shaking (side to side motion of paddle) 2. Hitting (up and down motion of paddle) 3. Pushing object

To copy an object from the object book onto the paddle the user simple places the paddle beside the desired object and the close proximity is detected and the object copied onto the paddle (figure 5a). Once a model is on the paddle it can be picked up and viewed from any viewpoint. To drop a model into the virtual scene the paddle is placed at the desired location and tilted until the model slides off (figure 5b). Models in the scene can be pushed around by pushing motions of the paddle (figure 6). A shaking motion is used to delete an object from the paddle.

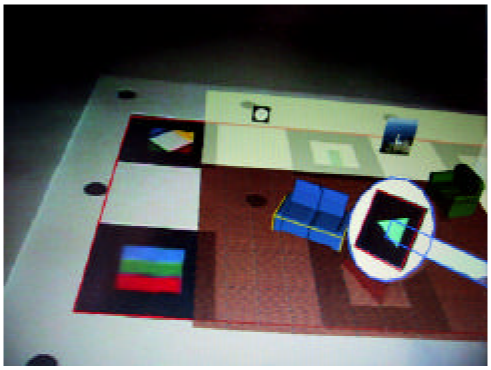

Figure 6 Moving virtual objects

As can be seen these interactions are very natural to perform with a real paddle, so in a matter of a few moments a user can assemble a fairly complex arrangement of virtual furniture. Using a single paddle and a variety of gestures we can assign multiple functions to create a time-multiplexed Tangible AR interfaces. This approach makes it very easy to interact with the application. Of course what the user is really doing is interacting with a simple CAD program, but instead of using a mouse or keyboard they are just manipulating a cardboard paddle in very intuitive ways.

3.3 The MagicBook: A Transitional Interface

The MagicBook project is an early attempt to explore how a physical object can be used to smoothly transport users between Reality and Virtuality. In this case a real book is used with tracking markers on each page. People can turn the pages of the book, look at the pictures, and read the text without any additional technology (figure 7a). However, if a person looks at the pages through an Augmented Reality display they see three-dimensional virtual models appearing out of the pages (figure 7b). The models appear attached to the real page so users can see the AR scene from any perspective simply by moving themselves or the book. The models can be any size and are also animated, so the AR view is an enhanced version of a three-dimensional "pop-up" book.

Users can change the virtual models simply by turning the book pages and when they see a scene they particularly like, they can fly into the page and experience the story as an immersive virtual environment (figure 7c). In the VR view they are free to move about the scene at will, so in the MagicBook people can experience the full Reality-Virtuality continuum.

Real books often serve as the focus for face to face collaboration and in a similar way the MagicBook interface can be used by multiple people at once. Several readers can look at the same book and share the story together. If they are using the AR displays they can each see the virtual models from their own viewpoint. Since they can see each other at the same time as the virtual models they can easily communicate using normal face to face communiation cues.

Multiple users can be immersed in the same virtual scene where they will see each other represented as virtual characters. More interestingly, one or more people may be immersed in the virtual world, while others are viewing the content as an AR scene. In this case those viewing the AR scene will see a miniature avatar of the immersive user in the virtual world (figure 8). In the immersive world, people viewing the AR scene appear as large virtual heads looking down from the sky. In this way people are always aware of where the other users of the interface are located and where their attention is focused.

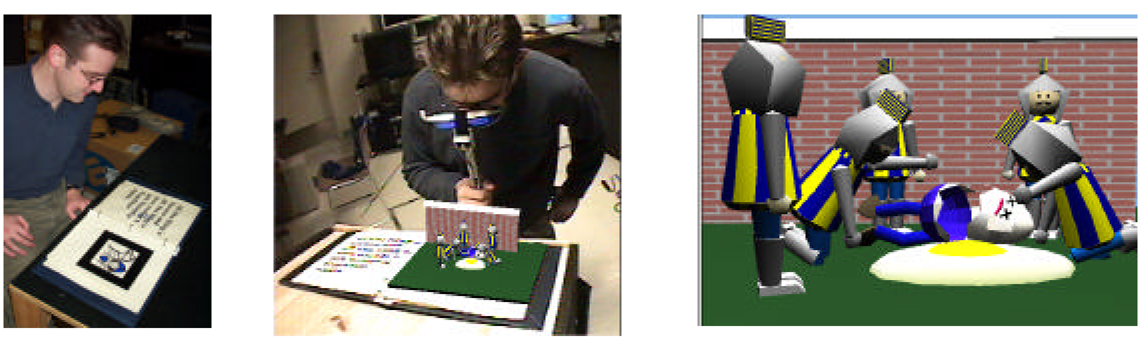

7a: Reality | 7b: Augmented Reality | 7c: Immersive Virtual Reality

Fig. 2: The MagicBook Transitional Interface

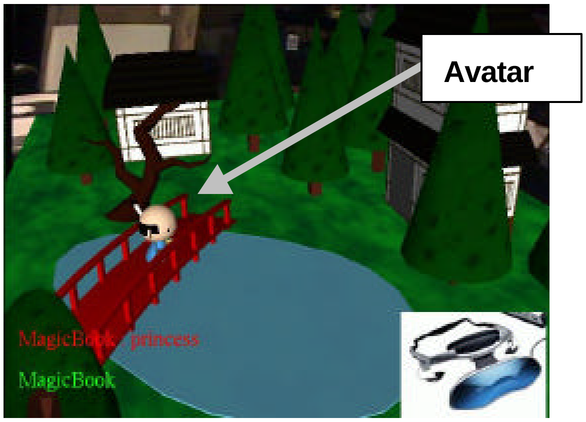

Figure 8: Avatar in an Exocentric AR view

Thus the MagicBook supports collaboration on three levels:

- *As a Physical Object:* Similar to using a normal book, multiple users can read the book together.
- *As an AR Object:* Users with AR displays can see virtual objects appearing on the pages of the book from their own viewpoint.
- *As an Immersive Virtual Space:* Users can fly into the virtual space together and see each other represented as virtual avatars.

The current MagicBook interface has two components; one or more a handheld displays (HHD) and the physical book. The HHD is a handle with a Sony Glasstron display mounted at the top, an InterSense InterTrax inertial tracker at the bottom, a small camera on the front of the Glasstron display and a switch and pressure pad (figure 9). The Sony Glasstron is a bioccular color display with two LCD panels of 265x235 pixel resolution. The camera output is connected to a desktop computer and the video-out of the computer is connected back into the HHD. So by looking through the HHD users experience a video-mediated reality.

Figure 9: The MagicBook Handheld Display

Users each have their own displays, so if two or more are looking at the same page, they will see the same virtual model attached to the page from their individual viewpoints. Since they can see each other at the same time they can use natural communication cues to enhance the collaboration.

When the user sees an AR scene they wish to explore, flicking the switch on the handle will fly them into an immersive VR environment. Head tracking is changed from the computer vision module to the InterTrax inertial orientation tracker so readers can look around the scene in any direction. By pushing the pressure pad on the handle they can fly in the direction they're looking. The harder they push the faster they fly. To return to the real world users simply need to flick the switch again.

When users are immersed in the virtual environment or are viewing the AR scenes their position and orientation is broadcast to the other users. This is used to place virtual avatars of people that are viewing the same scene, so users can collaboratively explore the virtual content.

The MagicBook software incoporates a complete VRML 97 parser making it easy for content developers to produce their own books. Nearly a dozen books have been created, including books for architecture, scientific visualization, education and entertainment. Users can dynamically load different book content simply by looking at the title pages.

4. Conclusions

Tangible Augmented Reality is a new approach to designing AR interfaces that emphasizes physical object form and interactions. It couples AR displays with tangible user interface controllers and 3D spatial interaction to enable a wide variety of powerful AR interfaces. Using design principles adapted from Tangible User Interfaces we can develop AR interfaces that support seamless interaction and are very intuitive to use. We believe that exploration with Tangible AR interfaces are a first step towards developing new physically-based interface metaphors that are unique to Augmented Reality.

There are several advantages of Tangible AR interfaces. First, Tangible AR interfaces are *transparent interfaces* that provide for seamless two-handed 3D interaction with both virtual and physical objects. They do not require participants to use or wear any special purpose input devices and tools, such as magnetic 3D trackers, to interact with virtual objects. Instead users can manipulate virtual objects using the same input devices they use in physical world – their own hands - leading to seamless interaction with digital and physical worlds. This property also allows the user to easily use both digital and conventional tools in the same working space.

Tangible AR allows *seamless spatial interaction* with virtual objects anywhere in their physical workspace. The user is not confined to a certain workspace but can pick up and manipulate virtual data anywhere just as real objects, as well as arrange them on any working surface, such as a table or whiteboard. The digital and physical workspaces are therefore continuous, naturally blending together.

Tangible AR interfaces can allow the design of a simple yet effective and *consistent AR interface model*, providing a set of basic tools and operations that allow users, for example, to add, remove, copy, duplicate and annotate virtual objects in AR environments.

An interesting property of Tangible AR interfaces is their *ad-hoc, highly re-configurable nature*. Unlike traditional GUI and 3D VR interfaces, Tangible AR interfaces are in some sense designed by user as they are carrying on with their work. In these interfaces the users are free to put interface elements anywhere they want: tables, whiteboards, in boxes and folders, arrange them in stacks or group them together. How the interface components should be designed for such environments, if they should be aware of the dynamic changes in their configuration, and how this can be achieved are interesting future research directions.

Another advantage is the use of *physical form-factor* to support interface functions. In our interfaces the physical design of the tangible interfaces elements provide affordances that suggest how they are to be used. Naturally the physical form factor and the computer graphics design of the virtual images attached to the interfaces is important and should correspond to each other.

Finally, Tangible AR interfaces *naturally support face-to-face collaboration*. People commonly use the resources of the physical world to establish a socially shared meaning. Physical objects support collaboration both by their appearance, the physical affordances they have, their use as semantic representations, their spatial relationships, and their ability to help focus attention. In a Tangible AR interface the physical objects can further be enhanced in ways not normally possible such as providing dynamic information overlay, private and public data display, context sensitive visual cues, and physically based interactions.

Acknowledgements

We would like to Ivan Poupyrev, Hirokazu Kato, and Desney Tan for their contributions to the interfaces shown. Without their work none of these interface would have been developed.

References

1. Billinghurst, M., Kato, H., Kraus, E., May, R. Shared Space: Collaborative Augmented Reality. In *Visual Proceedings, SIGGRAPH 99,* August 7-12th, Los Angeles, CA, ACM Press, 1999.

2. Billinghurst, M., Kato, H., Poupyrev, I. (2001) The MagicBook: Moving Seamlessly between Reality and Virtual Reality. IEEE Computer Graphics and Application, May 2001, Vol. 21, No. 3, pp. 2-4.

3. C. Elliot, G. Schechter, R. Yeung, and S. Abi-Ezzi: "TBAG: A High Level Framework for Interactive, Animated 3D Graphics Applications," *proceedings of ACM SIGGRAPH '94*, pp.421-434 (1994).

4. Fitzmaurice, G. and Buxton, W. (1997). An Empirical Evaluation of Graspable User Interfaces: towards specialized, space-multiplexed input. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'97). pp. 43-50. New York: ACM.

5. Fjeld, M., Voorhorst, F., Bichsel, M., Lauche, K., Rauter-berg, M., H., K. (1999). Exploring Brick-Based Navigation and Composition in an Augmented Reality. In *Proceedings of HUC' 99*, pp. 102-116.

6. Gorbet, M., Orth, M., Ishii, H. Triangles: Tangible Interface for Manipulation and Exploration of Digital Information Topography. In *Proceedings of CHI 98,* Los Angeles, CA, 1998.

7. Hedley, N., Postner, L., May, R., Billinghurst, M., Kato, H., (2001) Collaborative AR for Geographic Visualization. In proceedings of the International Symposium on Mixed Reality (ISMR 2001), March 14th-15th, Yokohama, Japan.

8. Ishii, H., Kobayashi, M., Arita, K. (1994). *Iterative design of seamless collaborative media.* CACM 37(8), 83-97.

9. Ishii, H., Ullmer, B. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In *Proceedings of CHI 97,* Atlanta, Georgia, USA, ACM Press, 1997, pp. 234-241.

10. Kato, H., Billinghurst, M. (1999) Marker Tracking and HMD Calibration for a video-based Augmented Reality Conferencing System. In *Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*. October, San Francisco, USA.

11. Kato, H., Billinghurst, M., Poupyrev, I., Imamoto, K., Tachibana, K. (2000) Virtual Object Manipulation on a Table-Top AR Environment. In proceedings of the International Symposium on Augmented Reality (ISAR 2000), Munich, Germany.

12. Poupyrev, I., Tan, D., Billinghurst, M., Kato, H., Regenbrecht, H., Tetsutani, N. (2001) Tiles: A Mixed Reality Authoring Interface. In Proceedings of INTERACT 2001, July 9-13, Yokohama, Japan.

13. Poupyrev, I., Berry, R., Billinghurst, M., Kato, H., Nakao, K., Baldwin, L., Kurumisawa, J. (2001) Augmented Reality Interface for Electronic Music Performance. In proceedings of the 9th International Conference on Human-Computer Interaction (HCI International 2001), August 5-10th, New Orleans, LA, USA.

14. Starner, T., Mann, S., Rhodes, B., Levine, J., Healey, J., Kirsch, D., Picard, R.W., Pentland, A. Augmented Reality through Wearable Computing. *Presence: Special Issue on Augmented Reality*, MIT Press, 1997.

15. Weiser, M. The Computer for the Twenty-First Century. *Scientific American*, 1991, 265 (3), pp. 94-104.

16. Wellner, P. Interactions with Paper on the DigitalDesk. *Communications of the ACM*, Vol. 36, no. 7, July 1993, pp. 87-96.