

The MagicBook: A Transitional AR Interface

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

Hirokazu Kato
Faculty of Information Sciences
Hiroshima City University, 3-4-1 Ozuka-Higashi,
Asaminami-ku, Hiroshima 731-3194, Japan
kato@sys.im.hiroshima-cu.ac.jp

Ivan Poupyrev
Interaction Laboratory
Sony CSL, 3-14-13 Higashi-Gotanda,
Tokyo 141-0022, Japan
poup@csl.sony.co.jp

Abstract

The MagicBook is a Mixed Reality interface that uses a real book to seamlessly transport users between Reality and Virtuality. A vision-based tracking method is used to overlay virtual models on real book pages, creating an Augmented Reality (AR) scene. When users see an AR scene they are interested in they can fly inside it and experience it as an immersive Virtual Reality (VR). The interface also supports multi-scale collaboration, allowing multiple users to experience the same virtual environment either from an egocentric or an exocentric perspective. In this paper we describe the MagicBook prototype, potential applications and user response.

Keywords: Augmented Reality, CSCW, Mixed Reality, Collaborative Virtual Environments.

1. Introduction

In the not too distant future an architect is meeting with her clients. On the table in their midst is a set of paper plans, sketches and renderings of what the final building design. Midway through the meeting the architect turns to her clients and says, “Let’s see what the building looks like with the landscaping you’ve chosen”. Each of the people around the table picks up a small handheld display and looks through it at the site plan. Superimposed on the paper map they see a three-dimensional virtual model of the building amidst trees and a flowing river. They can each see the model from their own perspective, moving around it to get a better look at the various landscaping features. Since the virtual model is overlaid on the real world they can also see each other and communicate naturally about the design changes they would like. After a few minutes the architect says, “Let’s look at the space inside the building”. With the flick of a switch the clients fly into the virtual model, becoming full immersed in a life-sized virtual reality scene. They are able to easily move through the foyer and into their building to inspect it from this new viewpoint. After exploring the new building, they can return to the real world just as easily to see the real building plan and discuss the changes they’d like made.

In this scenario the architect and her clients are able to move seamlessly between reality and immersive virtual reality while collaborating on the task at hand. Although it may seem far-fetched, this application could be demonstrated today with a new type of Mixed Reality interface that we have developed called the MagicBook. The MagicBook uses a physical object to move people effortlessly between the real and immersive virtual worlds while supporting natural collaboration. In this article we describe the fundamental ideas behind the MagicBook, its prototype implementation, possible application areas and some initial usability results.

2. The MagicBook Concept

Physicality, Augmented Reality and immersive Virtual Reality are traditionally separate media that people cannot seamlessly move between. In a Physical or Tangible User Interface (TUI) [12] real objects are used as the primary interface elements, in an AR interface virtual images are superimposed on the real world, while in a VR environment the real world is replaced entirely by computer generated imagery. In addition to single user applications, many computer interfaces have been developed that explore collaboration in a purely physical setting, in an AR setting, or in an immersive virtual world.

Our work involves exploring transitions between interface types, both in single user and collaborative applications. There are many possible ways of classifying computer interfaces. As Milgram points out [20], computer interfaces can be placed along a continuum according to how much of the users environment is computer generated (figure 1). On this Reality-Virtuality line Tangible User Interfaces lie far to the left, while computer-generated immersive virtual environments are placed at the rightmost extreme. In between are Augmented Reality interfaces where virtual imagery is added to the real world, and Augmented Virtuality interfaces where real world content is brought into immersive virtual scenes. Most current user interfaces can be placed at specific points along this line.




Figure 1: Milgram's Reality-Virtuality Continuum.

Similarly, Benford [4] classifies collaborative interfaces along two dimensions of Artificiality and Transportation. Transportation is the degree to which users leave their local space and enter into a remote space, and Artificiality the degree to which a space is synthetic or based on the physical world. Figure 2. shows the classification of typical collaborative interfaces. As can be seen, Milgram's continuum can be viewed as the equivalent of Benford's Artificiality dimension. Again, most collaborative interfaces exist at a discrete location in this two-dimensional taxonomy.




Figure 2: Benford's classification of Shared Space according to Transportation and Artificiality

However, human activity often cannot be broken into discrete components and for many tasks users may prefer to be able to easily switch between interfaces types, or co-located and remote collaboration. This is particularly true when viewing and interacting with three-dimensional graphical content. For example, even when using a traditional desktop modeling interface users will turn aside from the computer screen to sketch with pencil and paper. As Kiyokawa et. al. point out, immersive VR and AR are complimentary and the type of interface should be chosen according to the nature of the task [3]. For example if collaborators want to experience a virtual environment from different viewpoints or scale then immersive VR may be the best choice. However if the collaborators want to have a face-to-face discussion while viewing the virtual image an AR interface may be best. Similarly in a collaborative session, users may often want to switch between talking with their remote

collaborators and the people sitting next to them in the same location. Given that different degrees of immersion may be useful for different tasks and types of collaboration, an interesting question is how to support seamless transitions between the classification spaces.


The MagicBook project is an early attempt to explore how a physical object can be used to smoothly transport users between Reality and Virtuality, or between co-located and remote collaboration. Unlike other interfaces it cannot be placed as a discrete point on a taxonomy scale. In the remainder of this chapter we describe the MagicBook interface in more detail, the technology involved, initial user reaction and potential applications of the technology.

3. The MagicBook Experience

The MagicBook interface uses normal books as the main interface objects. People can turn the pages of these books, look at the pictures, and read the text without any additional technology (figure 3a). However, if they look at the pages through an Augmented Reality display they see three-dimensional virtual models appearing out of the pages (figure 3b). 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 traditional 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 3c). In the VR view they are free to move about the scene at will and interact with the characters in the story. Thus, users can experience the full Reality-Virtuality continuum.



3a: Reality



3b: Augmented Reality



3c: Immersive Virtual Reality

Figure 3: Using the MagicBook interface to move between Reality and Virtual Reality.

As can be seen the MagicBook interface has a number of important features:

1/ The MagicBook removes the discontinuity that has traditionally existed between the real and virtual worlds. VR is a very intuitive environment for viewing and interacting with computer graphics content, but in a head mounted display (HMD) a person is separated from the real world and their usual tools, or collaborators. The MagicBook enables people to move seamlessly between the real and virtual world and between co-located and remote collaborators.

2/ The MagicBook interface allows users to view graphical content from both egocentric and exocentric views, so they can select the viewpoint appropriate for the task at hand. For example an Augmented Reality viewpoint (exocentric view) may be perfect for viewing and talking about a model, but immersive VR (egocentric view) is better for experiencing the model at different scales or from different viewpoints. The MagicBook allows both AR and VR views to be used in a single interface.

3/ The computer interface has become invisible and the user can interact with graphical content as easily as reading a book. This is because the MagicBook interface metaphors are consistent with the form of the physical objects used. Turning a book page to change virtual scenes is as natural as rotating the page to see a different side of the virtual models. Holding up the AR display to the face to see an enhanced view is similar to using reading glasses or a magnifying lens. Rather than using a mouse and keyboard based interface users manipulate virtual models using real physical objects and natural motions. Although the graphical content is not real, it looks and behaves like a real object, increasing ease of use. This ease of interaction could lead to radically new forms of entertainment and educational applications usable by a wider group of consumers and members of the public.

3.1 Collaboration with the MagicBook

Physical objects, Augmented Reality interfaces and immersive VR experiences each have different advantages and disadvantages for supporting collaboration. As shown by Benford's classification, there has been a proliferation of collaborative interfaces, but it has traditionally been difficult to move between the shared spaces they create. For example, users in an immersive virtual environment are separated from the physical world and cannot collaborate with users in the real environment. The MagicBook supports all these types of interfaces and lets the user move smoothly between them depending on the task at hand.

Real objects 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 (fig. 4a). If these people then pick up their AR displays they will see the virtual models superimposed over the book pages from their own viewpoint. Since they can see each other and the real world at the same time as the virtual models they can easily communicate using normal face-to-face communication cues. All of the users using the MagicBook interface have their own independent view of the content so any number of people can view and interact with a virtual model as easily as they could with a real object (fig. 4b).



Figure 4a: Collaboration in the Real World





Figure 4b: Sharing an AR View

In this way the MagicBook technology moves virtual content from the screen into the real world, preserving the cues used in normal face-to-face conversation, and providing a more intuitive technology for collaboratively viewing three-dimensional virtual content.

Multiple users can also be immersed in the virtual scene where they will see each other represented as virtual characters in the story (fig. 5a). More interestingly, there may be situations where one or more users are immersed in the virtual world, while others are viewing the content as an Augmented Reality scene. In this case the AR user will see an exocentric view of a miniature figure of the immersed user, moving as they move themselves about the immersive world (figure 5b). Naturally, in the immersive world users viewing the AR scene appear as large virtual heads looking down from the sky. When users in the real world move, their virtual avatars move accordingly. In this way people are always aware of where the other users of the interface are located and where their attention is focused.



5a: Avatar in a Immersive Scene



5b: Avatar in an Exocentric AR view

Figure 5: Collaboration in the MagicBook

Thus the MagicBook interface supports collaboration on three levels:

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

The interface also supports collaboration on multiple scales. Users can fly inside the virtual scenes (an egocentric view) and see each other as virtual characters, but a non-immersed user will also see the immersed users as small virtual characters on the book pages (an exocentric view). This means that a group of collaborators can share both egocentric and exocentric views of the same game or data set, leading to enhanced understanding .

3.2. The MagicBook Interface

The MagicBook interface has three main components; a handheld Augmented Reality display (HHD), a computer, and one or more physical books. The books look like any ordinary book and have no technology embedded in them, while the display is designed to be easily held in one hand and to be as unencumbering as possible.



6a: The Hand Held Display



6b: A MagicBook Book

Figure 6: Components of the MagicBook Interface

Each user has their own handheld display and computer to generate an individual view of the scenes. These computers are networked together or exchanging information about avatar positions and the virtual scene each user is viewing. The handheld AR display is a handle with a Sony Glasstron PLM -A35 display mounted at the top, an InterSense InterTrax [11] inertial tracker at the bottom, a small color video camera on the front of the Glasstron display and a switch and pressure pad embedded in the handle (figure 7). The PLM-A35 is a low cost biocular display with two LCD panels of 260x230 pixel resolution.




Figure 7: The Handheld Display Components

The camera output is connected to the computer graphics workstation; computer graphics are overlaid on video of the real world and resultant composite image shown back in the Glasstron display. In this way users experience the real world as a video-mediated reality. One advantage of this is that the video frames that are being seen in the display are exactly the same frames as those drawn on by the graphics software. This means that the registration between the real and virtual objects appears almost perfect because there is no apparent lag in the system. The video of the real world is actually delayed until the system has completed rendering the 3D graphics. On a mid range PC (866Mhz Pentium III) with a virtual scene of less than 10,000 polygons we can maintain a refresh rate of 30 frames per second. This is fast enough that users perceive very little delay in the video of the real world and the virtual objects appear stuck to the real book pages.

Although commercially available hardware was used in this interface, the “Opera glass” form factor of the hand held display was deliberately designed to encourage seamless transition between Reality and Virtual Reality. Prior experience had taught us that many users are reluctant to wear head mounted displays, and that they are often difficult to get on and off. Users can look through the display to see AR and VR content, but can instantaneously return to viewing the real world simply by moving the display from in front of their eyes. The hand held display is far less obtrusive and easy to remove than any head worn display, encouraging people to freely transition along the Reality-Virtuality continuum. It is also easy to share, enabling several people to try a single display unit and see the same content as they pass it amongst themselves. The form factor of the handle itself went through many design iterations (see figure 8) to ensure that the components could be mounted on it, that it would be comfortable to hold, and that it could be easily passed around.




Figure 8. Alternative Handle Designs

The books used in the MagicBook interface are normal books with text and pictures on each page. Certain pictures have thick black borders surrounding them and are used as tracking marks for a computer vision based head tracking system. When the reader looks at these pictures through the HHD, computer vision techniques are used to precisely calculate the camera position and orientation relative to the tracking mark. The head tracking uses the ARToolKit tracking library, a freely available open-source software package for developing vision based AR applications [2]. Figure 9 summarizes how the ARToolKit tracking library works. Once the user's head position is known the workstation generates virtual images that appear precisely registered with the real pages. Our use of 2D markers for AR tracking is similar to the CyberCode work presented by Rekimoto [24] and other vision based tracking systems.




Figure 9: The ARToolKit Tracking Process

When the user sees an AR scene they wish to explore, flicking the switch on the handle will fly them smoothly into the scene, transitioning them into the immersive VR environment. In the VR scene users can no longer see the real world and so the head tracking is changed from the computer vision module to the InterTrax inertial orientation tracker. The output from the InterTrax inertial compass is used to set the head orientation in the virtual scene. The InterTrax provides three-degree of freedom orientation information with a high accuracy and very little latency. Readers can look around the scene in any direction and 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. The pressure pad and switch are both connected to a TNG interface box [29] that converts their output to a single RS-232 serial data signal.

The MagicBook application is also a networked application. Each of the user computers are networked together for exchanging information about avatar positions and the virtual scene that each user is viewing. When users are immersed in the virtual environment or are viewing the AR scenes their position and orientation is broadcast using TCP/IP code to a central server application. The server application then re-broadcasts this information to each of the networked computers and the MagicBook graphical client code. This is used to place virtual avatars of people that are viewing the same scene, so users can collaboratively explore the virtual content. Since each of the client applications contain a complete copy of the graphics code, only a very small amount of information needs to be exchanged. Thus MagicBook applications can potentially support dozens of users. There is also no need for users to be physically co-located. The virtual avatars can be controlled by users in the same location or remote from each other. So the MagicBook technology supports both face-to-face and remote collaboration.

3.2 MagicBook Applications

To encourage exploration in a number of different application areas we have developed the MagicBook as a generic platform that can be used to show almost any VRML content. VRML is a standard file format for three-dimensional computer graphics. We use an open source VRML rendering library called libVRML97 [21] that is based on the OpenGL low-level graphics library. Since VRML is exported by most 3D modeling packages, it is easy for content developers to build their own MagicBook applications. In addition, there are hundreds of thousands of VRML files freely available on the World Wide Web. Once the 3D content have been developed it is very simple to make the physical book pages and update the configuration files to load the correct content. In a matter of a few hours a new MagicBook application can be made.

This ease of development has resulted in the production of nearly a dozen books in a variety of application domains. Among others, we have a Japanese children's story that involves the reader in a treasure hunt, a macabre version of the Humpty Dumpty tale, a World War One History book, and a fanciful science fiction snowboard experience that allows the reader to ski Mt. St. Helens after it erupts. These MagicBook applications explore new literary ground where the reader can actually become part of the story and where the author must consider issues of interactivity and immersion.

The MagicBook technology has also strong application potential for scientific visualization. We have begun exploring using this technology for viewing geo-spatial models. Figure 10 shows views of typical oilfield seismic data superimposed over a tracking card. Currently petroleum companies deploy expensive projection screen based visualization centers around the world. The tracking systems used in the MagicBook interface are completely sourceless and so potentially mobile. In the near future it will be possible to run the MagicBook software from a laptop computer and so support a radically new way of presenting visualisation data in a field.




Figure 11: Seismic Data on a Tracking Marker

One of the more interesting applications is an educational textbook designed to teach architects how to build Gerrit Rietveld's famous Red and Blue Chair (figure 11). After a brief introduction to Rietveld's philosophy and construction techniques the readers are treated to a step-by-step instruction guide to building the chair. On each page is a two-dimensional picture of the current stage of the chair construction. When readers look at this page in their hand-held displays, they see a three-dimensional model of the partially completed chair popping out of page. On the final page they see a virtual model of the completed chair that they can fly into and see life-sized. Being able to see the chair from any angle during the construction process as well as a life-sized model at the end is a powerful teaching tool.




Figure 11: Stages in building Gerrit Rietveld's Red and Blue Chair

In a business setting each user could have their own independent view into the geo-spatial data set, and the number of users can be easily scaled to support both small and large group meetings. The opportunities for seeing private views of a shared dataset are particularly important when viewing sensitive data. The MagicBook gives engineers the ability to more naturally manipulate data sets and find obvious seismic features.

Medical education is another natural application area. A group of medical students could be sitting around a table, seeing each other as well as a three-dimensional animated model of the heart floating over a picture in a textbook. They could then fly into the model and see what it would look like from the perspective of someone walking through the four chambers. Medicine is an ideal application area for MagicBook technology because in current medical education textbooks and other 2D media to teach an inherently spatial subject.

Perhaps the most near term application of the MagicBook technology is for games and entertainment. Aside from interactive stories the MagicBook platform could be used to extend many already popular games. For example, by looking at a chess board through AR glasses users could see animated models of the various pieces; placing two pieces side by side could cause the characters to battle, and the user could fly into the battle scene to see it from a first person perspective.

3.4 User Feedback

"I think this is a great step towards immersive imagination."

"Great idea! I liked the handheld device."

"This is my favorite VR experience of the past few years - I like books and this "book extension" is really cool!"

"Amazing mixed reality, the motion tracking was excellent and there are many possible applications of the technology."

"You can turn the pages and put your head in there. It is very easy to change orientation by walking and turning your body. I loved the sense of control."

User Comments from Siggraph 2000

The MagicBook software was first shown at the Siggraph 2000 conference where over 2,500 people tried the books in the course of a week. Siggraph is a demanding environment to display an interactive experience because attendees typically only have a few minutes and need to be able to master the technology immediately. Although we did not have time for a rigorous user study, 55 of these people filled out a simple survey about their experience. The quotes above come from Siggraph attendees.

Feedback was very positive. People were able to use the interface with minimal training, they enjoyed the hand held displays, being able to view different AR scenes, and fly into the immersive VR worlds. Users felt that the interface was easy and intuitive to use. They were given the question “How easily could you move between the real and virtual worlds?”, and asked to respond on a scale of 1 to 7, where 1 was “not very easy” and 7 “very easy”. The average response was 5.9, and as the graph below (figure 12) shows users overwhelmingly felt they could easily transition between the real and virtual worlds.




Figure 12: How easy was it to move between Reality and Virtual Reality ? (7= very easy).

However users felt that it was not as easy to collaborate with each other (figure 13). On a scale of 1 to 7 with 1 being not very easy and 7 being very easy, the average response to the question “How easy was it to collaborate with others?” was 3.4. This was probably due to many of the people trying the books by themselves, or when using it with another person not being aware of their avatar in the scene.




Figure 13: How easy was it to Collaborate? (1 = not very easy, 7= very easy).

Another area for potential improvement is in the navigation metaphor used. We modelled movement in the immersive virtual world after movement in the real world and assumed that users would rarely want to move

backwards, since people rarely walk backwards. However it seems that users expected more of a video game metaphor and a majority of people immersed in the VR scenes asked how they could fly backwards. This is something to be addressed in future versions of the interface.

When asked what they liked most about the interface, users wrote that they most enjoyed the Augmented Reality scenes and the ability to fly into the virtual scenes. They were also impressed by the responsiveness of the camera tracking and the interactivity of the experience. The most common responses are listed below.

What did you like about the MagidBook exhibit?

- The Augmented Reality book scenes (8)
- How innovative and cool it was (8)
- The camera tracking (6)
- The idea (5)
- The ability to fly into the virtual scenes (4)
- The interactivity (3)

However users also felt that there were things that could be improved. They listed more realistic graphics, better animation and image quality as areas of concern. Many users also felt that movement through the virtual world could be improved, mostly because we did not allow users to move backwards. The most common responses are listed below.

What do you think could be improved in the MagicBook interface?

- The realism and complexity of the graphics content (16)
- Movement through the Virtual Space (8)
- The image quality (5)
- The latency in the graphics (4)
- The interactivity (3)

Users were also asked to write what applications they could imagine this technology being used for. Education (17 responses) and entertainment (11) were the most common responses, although people also thought this could be used as an interactive book (5), for training (5), visualization (3) and even teleconferencing (3). One particularly interesting suggested application was a reading comprehension system for children.

4. Related Work

The MagicBook project has arisen out of past work in Mixed Reality and CSCW interfaces developed by a wide variety of research groups. We are concerned with developing Augmented Reality technology to support face-to-face and remote collaboration, particularly for 3D CSCW. This section briefly reviews some of the prior research in the fields of remote and face-to-face collaboration that our work is built upon.

Early attempts at supporting face-to-face collaboration with computers were based around computer conference rooms. These were meeting rooms in which each of the participants sitting at a conference table had their own desktop computer and a video projector was used to provide a public display space. One example is the Colab room at Xerox [27] that was used to support small group face-to-face meetings. Colab used a network on workstations running distributed software applications designed to support brainstorming, document preparation, proposal evaluation and other collaborative processes. However there were very few successful computer conference rooms. One of the reasons for this is the discovery that collaborators work together better if they are focussed on a common workspace. So user's collaborating on separate workstations, even if they are side by side, do not perform as well as if they were huddled around a single machine [10].

Building on the notion of the common workspace, Weiser developed the concept of Ubiquitous Computing where computers disappeared into the physical environment and seamlessly supported collaboration through large common displays [30]. The importance of a central display for supporting face-to-face meetings has been recognized by the developers of large interactive displays (such as the LiveBoard [9]). Others have developed shared workspaces that allow physical and digital information to interact together in unique ways. For example, Wellner's DigitalDesk used a projector to display applications on a desktop surface and an overhead camera to recognize hand gestures and physical object placement to interact with the applications [31]. In this case the physical objects provide a common semantic representation as well as a tangible interface for the digital information space. Ishii's work on Tangible User Interfaces [12] has provided several examples of how physical objects can enhance both face-to-face and remote collaboration.

Although these types of applications are valuable for supporting collaboration on two-dimensional tasks, they have limited usefulness for three-dimensional CSCW. The earliest efforts in developing interfaces for 3D CSCW involved simply adding support for remote collaboration to existing screen-based single user 3D packages.

However a two-dimensional interface for three-dimensional collaboration can have severe limitations. For example, Li-Shu [18] developed a workstation based collaborative CAD package but users found it difficult to visualize the different viewpoints of the collaborators making communication difficult. Communication was also restricted to voice and pointing with a graphical icon, further compounding the problem.

Alternative techniques include using large parabolic stereo projection screens to project three-dimensional virtual images into space. CAVE-like systems [8] and the responsive workbench [16] allow a number of users to view stereoscopic 3D images by wearing LCD-shutter glasses. These images are projected on multiple large screen projection walls in the case of the CAVE, or a large opaque tabletop display for the responsive workbench. Unfortunately in both cases the images can be rendered from only a single user's viewpoint, so only one person will see true stereo. This makes it impossible for users to surround the Responsive Workbench table, or to spread themselves throughout the CAVE and see the correct stereoscopic image. The devices are also need bulky hardware such as a projection screen or large beam splitter, are not portable and require expensive optics.

Mechanical devices can also be used to create volumetric displays. These include scanning lasers onto a rotating helix to create a three-dimensional volumetric display [26] or using a rotating phosphor coated plate activated with electron guns [6]. These devices are also not portable and do not allow direct interaction with the images because of the rotating display surface.

In contrast immersive virtual reality has proven to be a natural medium for collaboration on spatial tasks. In this setting computers can provide the same type of collaborative information that people have in face-to-face interactions, such as communication by object manipulation and gesture [32]. Work on the DIVE project [7], GreenSpace [19] and other fully immersive multi-participant virtual environments has shown that collaborative work is indeed intuitive in such surroundings. Users can freely move through the space setting their own viewpoints and spatial relationships. Gesture, voice and graphical information can all be communicated seamlessly between the participants. Spatial visual and audio cues can combine in natural ways to aid communication [3]. However most current multi-user VR systems are fully immersive, separating the user from the real world. So even users who are co-located in the real world cannot use real non-verbal cues, or physical objects to aid their collaboration.

Augmented Reality technologies can be used to overcome this limitation. Amselen [1] and Rekimoto [23] have explored the use of tracked hand held LCD displays in collaborative environments. Amselen uses LCD panels as portable windows into a shared multi-user immersive environment, while Rekimoto attaches small cameras to LCD panels to allow virtual objects to be composited on video images of the real world and support face-to-face collaboration. In the Studierstube project of Schmalstieg et. al. [25] users wear see-through head mounted displays to collaboratively view 3D models of scientific data superimposed on the real world. They report users finding the interface very intuitive and conducive to real world collaboration because the groupware support can be kept simple and mostly left to social protocols. The AR² Hockey work of Ohshima et. al. [22] is also very similar. In this case two users wear see-through head mounted displays to play an Augmented Reality version of the classic game of air hockey. Like Schmalstieg, they report that users are able to naturally interact with each other and collaborate on a real world task. In our own work we have found that users perform better on a collaborative task in an AR interface than doing the same task in an immersive environment [5]. These results show that Augmented Reality interfaces allow multiple users to work in both the real and virtual world simultaneously facilitating CSCW in a seamless manner.

The MagicBook project combines elements of these previous tangible, AR and VR interfaces. However the key difference is the support for easy transitions between Reality and Virtuality and multi-scale collaboration. The ability of users to fly into miniature virtual worlds and experience them immersively was previously explored by Stoakley et. al. in the Worlds in Miniature (WIM) work [28]. They used miniature worlds to help users navigate and interact with immersive virtual environments at full-scale. Although we draw on the WIM metaphor, we expand on this work by supporting transitions between media forms of the Reality-Virtuality continuum, not just within the medium of immersive virtual reality itself. Another point of departure is that we focus on collaboration.

The WIM interface explored the use of multiple perspectives in a single user VR interface, while the CALVIN work of Leigh et. al. [17] introduced multiple perspectives in a collaborative VR environment. In CALVIN users could either be *Mortals* or *Deities* and view the VR world from either an egocentric or exocentric view respectively. CALVIN supported multi-scale collaborative between participants so that deities would appear like giants to mortals and vice versa. Deities also had an expanded set of interface controls that gave them more power over the interface. The MagicBook supports the multi-scale collaboration of CALVIN, but also adds an Augmented Reality component.

Perhaps most closely related to our work is the work of Kiyokawa et. al. [Kiyokawa 13,14,15]. They have developed a two-person shared AR interface for face-to-face computer-aided design. In their interface users can also change their body scale and experience the virtual world immersively, thus their work supports the seamless transition between an AR and immersive VR experience. As they point out, immersive VR is the best way for designers to get different perspectives on the virtual environment, while collaborative AR is the best way to support face-to-face discussion about the 3D design. Our work improves on their interface by adding a tangible user interface component and so supporting transitions along the entire Mixed Reality continuum, not just between AR and VR.

5. Conclusions

As computers become more ubiquitous and invisible there is a need for new interfaces that blur the line between Reality and Virtual Reality and let users easily move between the physical and digital domains. This can only be achieved by the use of Mixed Reality interfaces. The MagicBook is an early attempt at a transitional Mixed Reality interface for viewing and interacting with spatial datasets. In addition to allowing users to move between Reality and Virtual Reality at the flick of a switch, it also supports collaboration on multiple levels. Although the MagicBook supports viewing of sophisticated computer graphics content, the computer is invisible. Rather than using a mouse or keyboard, interaction is focused around a real book and a tangible interface. Even complete novices can use the interface and become part of the story.

Initial user feedback has been very positive and we are continuing to improve the interface. In the future we plan on exploring more intuitive ways for users to navigate through and interact with the virtual models. We are also working on ways of integrating the MagicBook approach into an environment with projective displays and so allow seamless transition between two-dimensional and three-dimensional views of a data set in a traditional office setting.

For more information about the MagicBook project and to download a free version of the ARToolKit software please visit <http://www.hitl.washington.edu/magicbook/>.

Acknowledgements

The authors would like to thank ATR MIC Labs for their continued support, Keiko Nakao, Susan Campbell, and Dace Campbell for making the models and books shown, and Tom Furness for creating a stimulating environment to work in.

References

- [1] D. Amselen. A Window on Shared Virtual Environments. *Presence*, 1995, Vol. 4(2), pp. 130-145.
- [2] ARToolKit website – http://www.hitl.washington.edu/research/shared_space/download/
- [3] S. Benford, S. and L. Fahlen. 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.
- [4] S. Benford, C. Greenhalg, G. Reynard, C. Briwn, B. Koleva. Understanding and Constructing Shared Spaces with Mixed Reality Boundaries. *ACM Transactions on Computer-Human Interaction (ToCHI)*, 5(3), pp. 185-223, September 1998, ACM Press.
- [5] M. Billinghurst, S. Weghorst, T. Furness. (1996) Shared Space: Collaborative Augmented Reality. *Workshop on Collaborative Virtual Environments 1996 (CVE 96)*, 19-20 Sept. 1996, Nottingham, UK.
- [6] B. Blundell, A. Schwarz. A Graphics Hierarchy for the Visualization of 3D Images by Means of a Volumetric Display System. In *Proceedings of the IEEE Region 10's Ninth Annual International Conference*, Singapore, Aug. 22-26, 1994, pp. 1-5. Vol. 1. IEEE New York, NY.
- [7] C. Carlson, O. Hagsand. DIVE - A Platform for Multi-User Virtual Environments. *Computers and Graphics*. Nov/Dec 1993, Vol. 17(6), pp. 663-669.
- [8] C. Cruz-Neira, D. Sandin, T. Defanti, R. Kenyon, and J. Hart. The CAVE: Audio Visual Experience Automatic Virtual Environment. *Communications of the ACM*, 1992, Vol. 35 (6), pp. 65.

- [9] S. Elrod, LiveBoard: A large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of CHI 92: Human Factors in Computing Systems*, Monterey, CA, 1992, pp. 599-607.
- [10] K. Inkpen. Adapting the Human Computer Interface to Support Collaborative Learning Environments for Children. PhD Dissertation, Dept. of Computer Science, University of British Columbia, 1997.
- [11] InterSense company website – <http://www.isense.com>
- [12] H. Ishii, B. Ullmer, Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In *proceedings of CHI 97*, Atlanta, Georgia, USA, ACM Press, 1997, pp. 234-241.
- [13] K. Kiyokawa, H. Iwasa, H. Takemura, N. Yokoya. Collaborative Immersive Workspace through a Shared Augmented Environment, In *Proceedings of the International Society for Optical Engineering '98 (SPIE '98)*, Vol.3517, pp.2-13, Boston, 1998.
- [14] K. Kiyokawa, H. Takemura, N. Yokoya. "SeamlessDesign: A Face-to-face Collaborative Virtual / Augmented Environment for Rapid Prototyping of Geometrically Constrained 3-D Objects," *Proceedings of the IEEE International Conference on Multimedia Computing and Systems '99 (ICMCS '99)*, Vol.2, pp.447-453, Florence, 1999.
- [15] K. Kiyokawa, H. Takemura, N. Yokoya. "A Collaboration Supporting Technique by Integrating a Shared Virtual Reality and a Shared Augmented Reality", *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC '99)*, Vol.VI, pp.48-53, Tokyo, 1999.
- [16] W. Kruger, C. Bohn, B. Frohlich, H. Schuth, W. Strauss, G. Wesche. The Responsive Workbench: A Virtual Work Environment. *IEEE Computer*, Vol. 28(7), 1995, pp.42-48.
- [17] J. Leigh, A. Johnson, C. Vasilakis, T. DeFanti, "Multi-perspective Collaborative Design in Persistent Networked Virtual Environments" In *Proceedings of IEEE Virtual Reality Annual International Symposium '96*. Santa Clara, California, Mar 30-Apr. 3, 1996, pp. 253-260.
- [18] Li-Shu, W. Flowers. Teledesign: Groupware User Experiments in Three-Dimensional Computer Aided Design. *Collaborative Computing*, Vol. 1(1), 1994, pp. 1-14.
- [19] J. Mandeville, J. Davidson, D. Campbell, A. Dahl, P. Schwartz, and T. Furness. A Shared Virtual Environment for Architectural Design Review. In *CVE '96 Workshop Proceedings*, 19-20th September 1996, Nottingham, Great Britain.
- [20] P. Milgram, F. Kishino, A Taxonomy of Mixed Reality Visual Displays. *IECE Trans. on Information and Systems (Special Issue on Networked Reality)*, vol. E77-D, no. 12, pp.1321-1329, 1994.
- [21] OpenVRML website – <http://www.openvrml.org/>
- [22] T. Ohshima, K. Sato, H. Yamamoto, H. Tamura. AR²Hockey: A case study of collaborative augmented reality, In *Proceedings of VRAIS'98*, 1998, IEEE Press: Los Alamitos, pp.268-295.
- [23] J. Rekimoto. Transvision: A Hand-held Augmented Reality System for Collaborative Design. In *Proceeding of Virtual Systems and Multimedia '96 (VSMM '96)*, Gifu, Japan, 18-20 Sept., 1996.
- [24] J. Rekimoto, Y. Ayatsuka. CyberCode: Designing Augmented Reality Environments with Visual Tags, Designing Augmented Reality Environments (DARE 2000), 2000.
- [25] D. Schmalstieg, A. Fuhrmann, Zs. Szalavári, and M. Gervautz. "Studierstube" – An Environment for Collaboration in Augmented Reality. *Proceedings of Collaborative Virtual Environments '96*, and *Virtual Reality Systems – Development and Applications*, vol. 3, no. 1, pp. 37-49, 1996.
- [26] P. Soltan, J. Trias, W. Dahlke, M. Lasher, M. McDonald, Laser-Based 3D Volumetric Display System: Second Generation. In *Interactive Technology and the New Paradigm for Technology*, IOP Press, 1995, pp. 349-358.

- [27] M. Stefik, G. Foster, D. Bobrow, K. Kahn, S. Lanning, L. Suchman, Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. In Communications of the ACM, January 1987, Vol 30, no. 1, pp. 32-47.
- [28] R. Stoakley, M. Conway, R. Pausch. Virtual Reality on a WIM: Interactive Worlds in Miniature. In Proceedings of CHI 95, 1995, ACM Press.
- [29] TNG 3B Interface Box available from Mindtel Ltd – <http://www.mindtel.com/>
- [30] Weiser, M. The Computer for the Twenty-First Century. Scientific American, pp. 94-104, 1991.
- [31] Wellner, P. Interactions with Paper on the DigitalDesk. Communications of the ACM, Vol. 36, no. 7, July 1993, pp. 87-96.
- [32] A. Wexelblat. The Reality of Cooperation: Virtual Reality and CSCW, in *Virtual Reality: Applications and Explorations*. Edited by A. Wexelblat. Boston, Academic Publishers, 1993.