Tiles: A Mixed Reality Authoring Interface

Ivan Poupyrev[1,i], Desney Tan[2,i], Mark Billinghurst[3], Hirokazu Kato[4, 6], Holger Regenbrecht[5] & Nobuji Tetsutani[6]

[1]Interaction Lab, Sony CSL
3-14-13 Higashi-Gotanda,
Tokyo 141-0022
Japan

[2] School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue,
Pittsburgh, PA 15213, USA

[3] University of Washington
[4] Hiroshima City University
[5] DaimlerChrysler AG
[6]ATR MIC Labs

poup@csl.sony.co.jp, desney@cs.cmu.edu, grof@hitl.washington.edu, kato@sys.im.hiroshima-cu.ac.jp, holger.regenbrecht@daimlerchrysler.com, tetutani@mic.atr.co.jp

Abstract: Mixed Reality (MR) aims to create user interfaces in which interactive virtual objects are overlaid on the physical environment, naturally blending with it in real time. In this paper we present *Tiles*, a MR authoring interface for easy and effective spatial composition, layout and arrangement of digital objects in MR environments. Based on a tangible MR interface approach, *Tiles* is a transparent user interface that allows users to seamlessly interact with both virtual and physical objects. It also introduces a consistent MR interface model, providing a set of tools that allows users to dynamically add, remove, copy, duplicate and annotate virtual objects anywhere in the 3D physical workspace. Although our interaction techniques are broadly applicable, we ground them in an application for rapid prototyping and evaluation of aircraft instrument panels. We also present informal user observations and a preliminary framework for further work.

Keywords: Augmented and mixed reality, 3D interfaces, tangible and physical interfaces, authoring tools

1 Introduction

Virtual objects are pervading our living and working environments, augmenting and even replacing physical objects. Electronic billboards are starting to replace familiar paper billboards in public spaces; and signs providing directions are often projected, rather then made out of the physical plastic or paper.

Mixed Reality research takes this integration between physical and virtual worlds even further. MR systems create advanced user interfaces and environments where interactive virtual objects are overlaid on the 3D physical environment, naturally blending with it in real time (Azuma, 1997; Milgram, Takemura, Utsumi, *et al.*, 1994). There are many potential uses for such interfaces, ranging from industrial, to medical and entertainment applications (e.g. Bajura, Fuchs et al. 1992; Poupyrev, Berry et al. 2000, see also Azuma, 1997 for survey).

In our work, we are interested in applying MR techniques to the task of collaborative design (Fjeld, Voorhorst, Bichsel, *et al.*, 1999; Kato, Billinghurst, Poupyrev, *et al.*, 2000). In one scenario, several architects and city planners gather around a conventional physical model of the city to evaluate how proposed buildings would alter the city appearance. Instead of using physical models of new buildings, the participants manipulate virtual 3D graphics models that are correctly registered and superimposed on the physical city model. The new buildings are virtual, so they can be quickly altered on the fly, allowing designers to evaluate the alternatives and possible solutions. Dynamic simulations, such as traffic flow and pollution can be simulated and superimposed right on the physical city model.

Unlike virtual reality (VR) interfaces, MR do not remove users from their physical environment. Users still have access to conventional tools and information, maps, and design schemes. Users can also continue to see each other and use gestures or facial expressions to facilitate their communication and enhance the decision process. Furthermore, as they proceed with their discussion they are implicitly documenting the design process by marking and annotating both virtual and physical objects.

This scenario remains mostly hypothetical. Most current MR interfaces work as information browsers allowing users to see virtual information embedded

[i] This work was conducted while the author was working at the ATR MIC Labs, Japan

into the physical world. However, few provide tools that let the user interact, request or modify this information effectively and in real time (Rekimoto, et al. 1998). Even the basic interaction tasks and techniques, such as manipulation, coping, annotating, dynamically adding and deleting virtual objects to the MR environment have been poorly addressed.

The current paper presents *Tiles*, a MR authoring interface that investigates interaction techniques for easy and effective spatial composition, layout and arrangement of digital objects in mixed reality environments. Several features distinguish *Tiles* from previous work. First, *Tiles* is a transparent interface that allows seamless two-handed 3D interaction with both virtual and physical objects. *Tiles* does not require participants to use or wear any special purpose input devices, e.g. 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. Second, unlike popular table-top based AR interfaces, where the virtual objects are projected on and limited by the 2D surface of a table (e.g. Rekimoto and Saitoh, 1999), *Tiles* allows full 3D spatial interaction with virtual objects anywhere in their physical workspace. The user can pick up and manipulate virtual data just as real objects, as well as arrange them on any working surface, such as a table or whiteboard. Third, *Tiles* allows the user to use both digital and physical annotations of virtual objects, using conventional tools such as PostIt™ notes. Finally, in *Tiles* we attempt to design a simple yet effective interface for authoring MR environments, based on a consistent interface model, providing a set of tools that allows users to add, remove, copy, duplicate and annotate virtual objects in MR environments. Although 2D and 3D authoring environments have been one of the most intensively explored topics in desktop and VR interfaces (e.g. Butterworth, Davidson, Hench, *et al.*, 1992; Mapes and Moshell, 1995) there are far fewer attempts to develop authoring interfaces for mixed reality. We discuss some of them in the next section.

2 Related work

We spend a significant part of our everyday life arranging and assembling physical objects in our workspace: books, papers, notes and tools. In recent years there has been a trend towards developing computer interfaces that also use physical, tangible objects for input devices. For example, in the Digital Desk project (Wellner, 1993), the position of paper documents and the user's hands on an augmented table were tracked using computer vision techniques. In this system, the user could seamlessly arrange and annotate both real paper and virtual documents using the same physical tool – a conventional pen. This approach was extended with graspable and tangible interfaces, which have been proposed as a possible interface model for such environments. This idea suggests using simple physical objects tracked on the surface of a table as either physical handles allowing to select, translate and rotate electronic objects or as data transport devices (Fitzmaurice, Ishii and Buxton, 1995; Fjeld, *et al.*, 1999; Ishii and Ullmer , 1997; Ullmer and Ishii, 1997; Ullmer, Ishii and Glas, 1998). Alternatively, Rekimoto, et al. (1999) used a special purpose laser pointer device and Hyperdragging interaction technique to move electronic documents between the computer and a shared workspace.

The main advantage of this approach is that the user does not have to wear any special-purpose display devices, such as a head-mounted display (HMD). Furthermore, physical, tangible interfaces allow the user to seamlessly interact with both electronic and physical objects simply with hands and physical tools, e.g. pen and wood blocks. However, because the output is limited to the 2D surface of the table, the user is not able pick up virtual documents and manipulate them freely in space as can be done with real paper documents. This interaction is also limited to flat paper-like objects. Presentation and manipulation of 3D virtual objects in such environments, though possible, is difficult and inefficient (Fjeld, *et al.*, 1999). Hence, these interfaces introduce *spatial seams* [i] in mixed reality environments – the interfaces are localized on an augmented surface and cannot extend beyond it.

Another fundamental alternative approach to building mixed reality workplaces is three-dimensional Augmented Reality (AR) (Azuma, 1997). In this approach, virtual objects are registered in 3D physical environments using magnetic or computer vision tracking techniques and then presented to the user looking through a HMD (e.g. Bajura, *et al.*, 1992; Feiner, MacIntyre and Seligmann, 1993) or a handheld display device (e.g. Fitzmaurice, 1993; Rekimoto and Nagao, 1995). Unlike tabletop-based MR, this approach allows the system to render 3D virtual objects anywhere in the physical environment to provide spatially seamless MR workspaces.

However, as Ishii points out, most AR researchers are primarily concerned with "considering purely visual augmentations" rather than the interaction and physical context of AR environments (Ishii and Ullmer , 1997). This has led to difficulty with designing interaction techniques that would let the user effectively manipulate 3D virtual objects distributed freely in a 3D workspace. Previous approaches to solve this problem include using a special purpose 3D input device to select and manipulate virtual ob-

[i] 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 (Ishii, Kobayashi and Arita, 1994).

jects, such as magnetic trackers used in Studierstube (Schmalsteig, Fuhrmann, Szalavari, *et al.*, 1996) and MARS systems (Hollerer et al. 1999). Traditional input devices, such as a hand-held mouse or tablet (Hollerer, *et al.*, 1999; Rekimoto, *et al.*, 1998), as well as speech input and intelligent agents (Anabuki, Kakuta, Yamamoto, *et al.*, 2000) have also been investigated. The major disadvantage with these approaches is that the user is forced to use two different interfaces – one for the physical and one for the virtual objects. Thus, the natural workflow is broken with *interaction seams* – every time the user needs to manipulate virtual objects, he or she needs to use a special purpose input device that would not be normally used in real world interaction.

Thus the current design of mixed reality interfaces, falls into two orthogonal approaches: tangible interfaces and tabletop MR offer seamless interaction but results in spatial discontinuities, while 3D AR provides spatially seamless mixed reality workspaces but introduces discontinuities in interaction. This paper presents an approach that merges the best qualities of both interaction styles. The *Tiles* system was developed to provide true spatial registration and presentation of 3D virtual objects anywhere in the physical environment. At the same time we implement a tangible interface that allows users to interact with 3D virtual objects without using any special purpose input devices. Since this approach combines tangible interaction with AR display we refer to it as *Tangible Augmented Reality*. In the next section we show how the *Tangible AR* can be used to build a simple yet effective MR authoring interface.

3 *Tiles* Interface

Tiles is a collaborative *Tangible AR* interface that allows several participants to dynamically layout and arrange virtual objects in a mixed reality workspace. In this system, the user wears a light-weight head-mounted display (HMD) with a small camera attached, both of which are connected to a computer. Output from the camera is captured by the computer which then overlays virtual images onto the video in real time. The resulting augmented view of the real world is then presented back to the user on his or her HMD so the user sees virtual objects embedded in the physical workspace (Figure 1 and Figure 2). The 3D position and orientation of virtual objects is determined using computer vision tracking techniques, tracking 3D position and orientation of square fiduciary markers that can be attached to any physical object. The tracking techniques have been inspired by Rekimoto (1988) and are more completely described in (Kato and Billinghurst, 1999) The virtual objects are rendered relative to these markers, and by manipulating marked physical objects, the user can manipulate virtual objects without need to use any additional input devices.

The rest of this section presents the *Tiles* interface and interaction techniques. Although our interface techniques are broadly applicable, the *Tiles* system has been developed for rapid prototyping and evaluation of aircraft instrument panels, a joint research initiative carried out with support from DASA/EADS Airbus and DaimlerChrysler AG. To ground further discussion and illustrate the rationale for our design decisions, we present a brief overview of the application design requirements.

3.1 Design Requirements

The design of aircraft instrument panels is an important procedure that requires the collaborative efforts of engineers, human factor specialists, electronics designers, airplane pilots and many others. Because mistakes are normally detrimental to aircraft safety, designers and engineers are always looking for new technologies that can reduce the cost of designing, prototyping, and evaluating the instrumental panels without compromising design quality. Since they are often building upon existing functional instruments, designers have taken a special interest in MR interfaces. This is because they often need to evaluate prototypes of instruments relative to existing instrumental panels, without having to physically build them. This design activity is inherently collaborative and involves team-based problem solving, discussions and joint evaluation. It also involves heavy use of existing physical plans, documents and tools.

Using observations of how instrument panels are currently designed, DASA/EADS Airbus and DaimlerChrysler engineers produced a set of requirements for MR interfaces to support this task. They envisioned MR interfaces allowing groups of designers, engineers, human factors specialists, and aircraft pilots to collaboratively outline and layout a set of virtual aircraft instruments on a board simulating an airplane cockpit. Designers would need to be able to easily add and remove virtual instruments from the board using a catalog of the virtual instruments. After the instruments are placed on the board, they would like to evaluate and rearrange the position of the instruments as necessary. The interface should also allow the use of existing physical schemes and documents with conventional tools, e.g. whiteboard markers, to let participants document solutions and problems, as well as add physical annotations to virtual instruments. A further requirement was that the resulting interface be intuitive, easy to learn and use.

3.2 Interface

3.2.1 Basics: Tiles interface components

The *Tiles* workspace and interface consist of: 1) a metal whiteboard in front of the user; 2) a set of

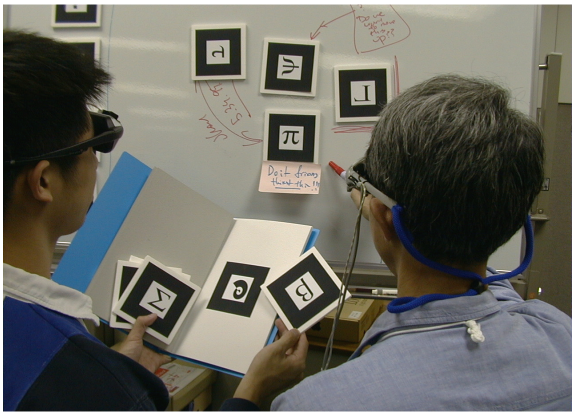

Figure 1: *Tiles* environment: users collaboratively arrange data on the whiteboard, using tangible data containers, *data tiles*, as well as adding notes and annotations using traditional tools: whiteboard pen and notes.

Figure 2: The user, wearing lightweight head-mounted display with mounted camera, can see both virtual images registered on tiles and real objects.

cardboard cards (15 by 15 centimetres each) with tracking patterns attached to them, which we call *tiles*. Each of these cards has a magnet on the back so it can be placed on and removed from the whiteboard; 3) a book, with marked pages, which we call *book tiles,* and 4) conventional tools used in discussion and collaboration, such as whiteboard pens and PostIt™ notes (Figure 1 and Figure 2).

The whiteboard acts as a shared collaborative workspace, where users can rapidly draw rough layout of virtual instruments using whiteboard markers, and then visualize this layout by placing and arranging tiles with virtual instruments on the board.

The tiles act as generic tangible interface controls, similar to icons in a GUI interface. So instead of interacting with digital data by manipulating icons with a mouse, the user interacts with digital data by physically manipulating the corresponding tiles. Although the tiles are similar to physical icons (phicons), introduced in metaDesk system (Ullmer and Ishii, 1997), there are important differences. In metaDesk, the authors proposed a close coupling between physical properties of phicons, i.e. their shape and appearance, to virtual object that phicons represent. For example, the shape of phicons representing a certain building had an exact shape of that particular building. In designing the *Tiles* interface we attempted to decouple physical properties of tiles from the virtual data as much as possible – the goal was to design universal data containers that can hold *any* digital data or no data at all. Interaction techniques for performing basic operations such as putting data on tiles and removing data from tiles are the same for all tiles, resulting in a consistent and streamlined user interface. This is not unlike GUI interfaces, where all basic operations on icons are the same irrespective of whether they represent a document or a game program – i.e. the user can move, open, resize and delete icons. Furthermore, because the user can dynamically put any digital data on the tile, our system does not require an excessive number of tiles, since they can be "recycled".

3.2.2 Classes of tiles: data, operators and menu

Not all tiles are the same – we use three classes of tiles: *data tiles, operator tiles* and *menu tiles*. All tiles share similar physical appearances and common operation. The only difference in their physical appearance is the icons identifying tile types. This allows users who are not wearing a HMD to identify the tiles purpose. Below we briefly summarize the basic properties of each of the classes:

- *Data tiles* are generic data containers. The user can put and remove virtual objects from the data tiles; if a data tile is empty, nothing is rendered on it. We use Greek symbols as tracking patterns to identify the data tiles.
- *Operator tiles* are used to perform basic operations on data tiles. Currently implemented operations include *deleting* a virtual object from a data tile, *copying a* virtual object to the clipboard or from clipboard to the data tile, and requesting *help* or annotations associated with a virtual object on the data tile. Iconic patterns are used to identify each operator tile, for example the tile that deletes a virtual object from data tiles is identified with a trashcan icon. In MR the operator tiles are also identified by virtual 3D widgets attached to them.
- *Menu tiles* make up a book with tiles attached to each page (Figure 1). This book works like a catalogue or a menu: as the user flips through the pages, he can see virtual objects attached to each page, choose the required instrument and then copy it from the book to any empty data tile.

3.2.3 Operations on tiles

All tiles can be manipulated in space and arranged on the whiteboard: the user simply picks up any of the tiles, examines its contents and places it on the whiteboard. Operations *between tiles* are invoked by bringing two tiles next to each other (within a distance less then 15% of the tile size).

For example, to copy an instrument to the data tile, the user first finds the desired virtual instrument in the menu book and then places any empty data tile next to the instrument (Figure 7). After a one second delay to prevent an accidental copying, a copy of the instrument smoothly slides from the menu page to the tile and is ready to be arranged on the whiteboard. Similarly, if the user wants to "clean" data from tile, the user brings the trashcan tile close to the data tiles, removing the instrument from it (Figure 3).

Using the same technique we can implement copy and paste operations using the *clipboard operator*: the user can copy an instrument from any of the data tiles to the clipboard and then from clipboard to an empty data tile (Figure 4). The current content of the clipboard is always visible on the virtual clipboard icon. There can be as many clipboards as needed – in the current implementation we have two independent clipboards.

Figure 3: The user cleans data tiles using trash can operator tile. The removed virtual instrument is animated to provide the user with smooth feedback.

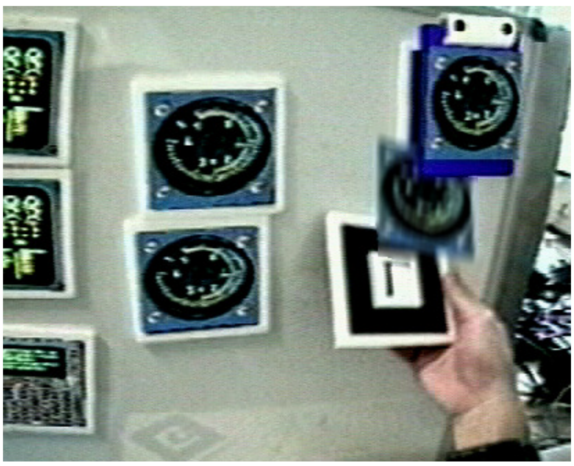

Figure 4: Coping data from clipboard to an empty data tile.

Table 1 summarises the allowed operations between tiles. Note that we have not defined any operations between data tiles because this would cause interaction between data tiles and not allow the user to lay them next to each other on the whiteboard.

3.2.4 Getting help in Tiles

Help systems have been one of the corner stones in providing guidance to users in a GUI, and effective MR interfaces will also require effective on-line help facilities. Therefore, we implemented a help tile: to receive help on any virtual object, the user simply places the help tile next to the data tile on which they require help. In the simplest case, this triggers explanatory text that appears within a bubble next to the help icon (Figure 5). Currently, this function is used by the designer to leave short digital annotations on the virtual instruments and to provide help for users while they manipulate the operator tiles.

3.2.5 Mixing physical and virtual tools in Tiles

The *Tiles* interface allows the users to seamlessly combine use of conventional physical tools, such as whiteboard pens, together with the virtual tools that we introduced in the previous sections. For example, the user can physically annotate a virtual aircraft instruments using a standard whiteboard pen or sticky note (see Figure 1, 2 and 6).

3.2.6 Collaboration

Tiles has been designed with collaboration in mind and allows several users interact in a same augmented workspace. We have been evaluating two possible scenarios: 1) All users are equipped with HMDs and can directly interact with virtual objects (Figure 1) and 2) Non-immersed users, i.e. users that do not wear HMDs collaborate with immersed users using an additional monitor presenting the view of immersed collaborator (Figure 7).

2.1 Initial User Feedback

Although the *Tiles* system has not yet been evaluated in rigorous user studies we have presented the interface in several public settings and received informal feedback from typical users. The *Tiles* system was first demonstrated at the IEEE/ACM International Symposium for Augmented Reality (ISAR) 2000 in Munich, Germany. About seventy users tested the system. We observed that with simple instructions, most of these users were able to quite effectively simulate the design process, laying out and rearranging the instruments on the board. They found the system easy to use, intuitive and quite enjoyable. DaimlerChrysler design engineers found that the concept meets the basic requirement for the authoring of MR environments and thought it promising enough to start evaluating its feasibility in real industrial applications.

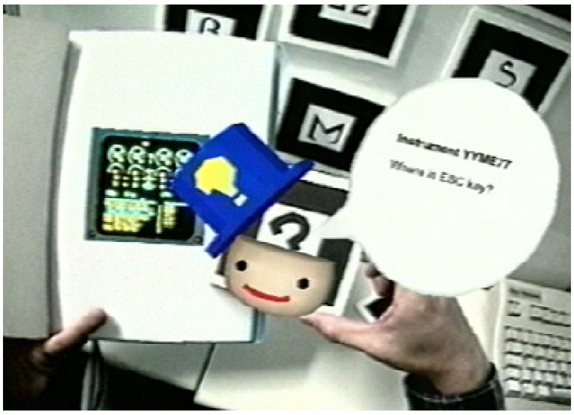

Figure 5: The user invokes an electronic annotation attached to the virtual objects using the help tile

Figure 6: Physically annotating virtual objects in *Tiles*

Figure 7: Collaboration between immersed and non-immersed users in *Tiles* environment

The most prevalent complaint was the physical design of the tiles. In designing the system, we wanted to keep the physical tiles as small as possible so as to match the size of the actual instruments. However, we tried to make the markers large enough for reliable tracking. As a result, the border around the tracked area, on which the user could place their fingers when holding the card, was uncomfortably small. Furthermore, the users tended to occlude the tracking border, which resulted in tracking failure. We are currently exploring different physical designs for the tiles in the next version of the system.

Our initial experiments with the non-immersed collaboration mode was encouraging in that the users were able to collaborate rather effectively. All interface components are simple physical objects identified with graphical icons, so the non-immersed user was able to perform the same authoring tasks as immersed user, i.e. laying out the tiles on the whiteboard, evaluating it, copying the virtual instruments on the data tiles and etc. We are planning to perform more extensive studies of this collaboration mode.

2.2 Implementation

The fundamental elements of any MR systems are techniques for tracking user position and/or viewpoint direction, registering virtual objects relative to the physical environment, rendering, and presenting them to the user.

The *Tiles* system is implemented using ARToolKit, a custom video see-through tracking and registering library (Kato and Billinghurst, 1999). We mark 15x15 cm paper cards with simple square fiduciary patterns consisting of thick black border and unique symbols in the middle identifying the pattern. The system does not have restrictions on symbols used for identification as long as it is asymmetrical to distinguish between the 4 possible orientations of the square border. The user wears a Sony Glasstron PLMS700 headset, which is lightweight and comfortable and provides VGA 800 by 600 pixel resolution. This was sufficient for reading text images rendered in our MR environment. A miniature NTSC Toshiba camera with a wide-angle lens (2.2 mm) is attached to the headset. The video stream from the camera is captured at 640x240 resolution to avoid interlacing problems and scaled back to 640x480 by using a line doubling technique.

After the computer vision pattern tracking identifies localization marks in the video stream, the relative position and orientation of the marks relative to the head-mounted camera can be determined and virtual objects can then be correctly rendered on top of the physical cards. Although the wide-angle lens distorts the video image, our tracking techniques are robust against these distortions and able to correctly track patterns without losing performance.

All virtual objects are represented as VRML97 models and a custom VRML browser has been built to manipulate and render 3D objects into the video stream. In the current *Tiles* application the system tracks and recognize 21 cards in total. The software is running on an 800Mhz Pentium III PC with 256Mb RAM and the Linux OS. This produces tracking and display rate of between 25 and 30 frames per second.

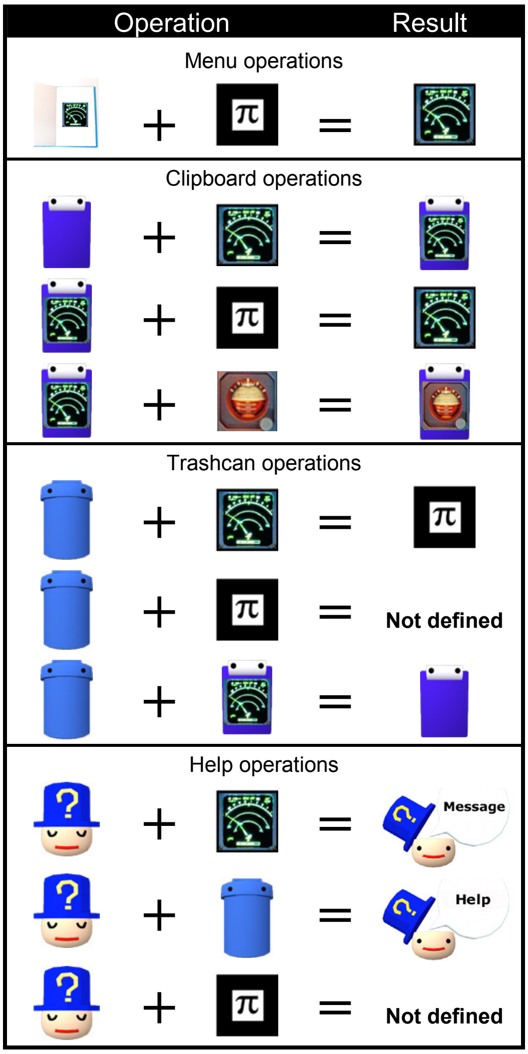

Table 1: Operations defined for different tiles types: e.g. bringing together menu tile and empty data tile will move instrument on the tile (first row in the table).

4 Discussion and Future Work

The *Tiles* system is a prototype tangible augmented reality authoring interface that allows a user to quickly layout virtual objects in a shared workspace and easily manipulate them without need of special purpose input devices. We are not aware of any previous interfaces that share these properties. In this section we discuss some of the *Tiles* design issues and future research directions.

Generality of Tiles, other applications. The interface model and interaction techniques introduced in *Tiles* can be easily extended to other applications that require mixed reality interfaces. Object modification techniques, for example, can be quite easily introduced into *Tiles* by developing additional operator cards that would let the user dynamically modify objects, e.g. scale them, change their colour and so on. We are also currently exploring more direct techniques that would track users' hands and allow the user to touch and scale virtual objects directly with gestures.

Although developing additional interaction techniques would allow *Tiles* to be used in many different application scenarios, we should note that in MR environments the user can easily transfer between the MR workspace and a traditional environments such as a desktop computer. Therefore, we believe that the goal of developing MR interfaces is not to bring every possible interaction tool and technique into the MR workspace, but to balance and distribute the features between the MR interface and other media: some tools and techniques are better for MR, some are better to be left for traditional tools. Hybrid mixed reality interfaces have been suggested by a number of researchers and are an interesting and important research direction (Schmalstieg, Fuhrmann and Hesina, 2000)

Ad-hoc, re-configurable interfaces. An interesting property of mixed reality interfaces is their ad-hoc, highly re-configurable nature. Unlike the traditional GUI and 3D VR interfaces, where the interface layout is mostly determined by an interface designer in advance, the MR interfaces are in some sense designed by user as they are carrying on with their work. Indeed, in *Tiles* 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 research directions.

Physical form-factor. Our initial user observations showed that in designing tangible MR interfaces, the form factor becomes an important design issue. Indeed, the main problem reported with *Tiles* was that the cards were too small, so people tended to occlude the tracking markers. In MR interfaces both the physical design of the interfaces and the computer graphics design of virtual icons attached to the interfaces is important. The design of physical components can convey additional semantics of the interface, for example the shape of the physical cards can be designed so that they can snap into each other as pieces in a jigsaw puzzle, and depending on their physical configuration resulting functionality of the interface could be different. Expressing different interface semantics by explicitly using the shape of the interface components can also be explored further in *Tiles* environment.

Remote and face-to-face collaboration. The current *Tiles* interface provides only very basic collaborative capabilities for co-located users. We are planning to explore remote collaboration techniques in the *Tiles* interface by using a digital whiteboard and global static camera to capture the writings on the whiteboard and location of tiles, and then distribute this to remote participants.

5 Conclusions

In this paper we presented *Tiles*, a MR authoring interface for easy and effective spatial composition, layout and arrangement of digital objects in MR environments. Based on a tangible MR interface approach, *Tiles* is a transparent user interface that allows users to seamlessly interact with both virtual and physical objects and introduces a consistent MR interface model, providing users a set of tools that allow dynamically to add, remove, copy, duplicate and annotate virtual objects anywhere in the 3D physical workspace. Although our interaction techniques are broadly applicable, we grounded them in an application for rapid prototyping and evaluation of aircraft instrument panels, a joint research initiative carried out with support from DASA/EADS Airbus. Informal user observations were encouraging and a framework for further work has been outlined.

References

Anabuki, M., Kakuta, H., Yamamoto, H., Tamura, H. (2000). Welbo: An Embodied Conversational Agent Living in Mixed Reality Spaces. *In Proceedings of the CHI'2000, Extended Abstracts*, ACM, pp. 10-11.

Azuma, R. (1997). A Survey of Augmented Reality. *Presence,* 6(4), pp. 355-385.

Bajura, M., Fuchs, H., Ohbuchi, R. (1992). Merging Virtual Objects with the Real World: Seeing Ultrasound Imagery Within the Patient. *In Proceedings of the SIGGRAPH'92*, ACM, pp. 203-210.

Butterworth, J., Davidson, A., Hench, S., Olano, T. (1992). 3DM: a three dimensional modeler using a head-mounted display. *In Proceedings of the Symposium on Interactive 3D graphics*, ACM, pp. 135-138.

Feiner, S., MacIntyre, B., Seligmann, D. (1993). Knowledge-Based Augmented Reality. *Communications of the ACM,* 36(7), 53-62.

Fitzmaurice, G., Ishii, H., Buxton, W. (1995). Bricks: Laying the foundations for graspable user interfaces. *In Proceedings of the CHI'95*, ACM, pp. 442-449.

Fitzmaurice, G. W. (1993). Situated information spaces and spatially aware palmtop computers. *Communication of the ACM,* 36(7), 38-49.

Fjeld, M., Voorhorst, F., Bichsel, M., Lauche, K., Rauterberg, M. et al. (1999). Exploring Brick-Based Navigation and Composition in an Augmented Reality. *In Proceedings of the HUC99*, pp. 102-116.

Hollerer, T., Feiner, S., Terauchi, T., Rashid, G., et al. (1999). Exploring MARS: developing indoor and outdoor user interfaces to a mobile augmented reality system. *Computers & Graphics,* 23, pp. 779-785.

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

Ishii, H., Ullmer , B. (1997). Tangible bits towards seamless interfaces between people, bits and atoms. *In Proceedings of the CHI'97*, ACM pp. 234-241

Kato, H., Billinghurst, M. (1999). Marker Tracking and HMD Calibration for a Video-based AR Conferencing System, *2nd Int. Workshop on AR*, pp.85-94.

Kato, H., Billinghurst, M., Poupyrev, I., Imamoto, K., Tachibana, K. (2000). Virtual Object Manipulation on a Table-Top AR Environment, *ISAR*, pp. 111-119

Mapes, D., Moshell, J. (1995). A Two-Handed Interface for Object Manipulation in Virtual Environments. *Presence*, 4(4), pp. 403-416.

Milgram, P., Takemura, H., Utsumi, A., Kishino, F. (1994). Augmented Reality: A Class of Displays on the Reality-Virtuality Continuum, *SPIE*, 2351

Poupyrev, I., Berry, R., Kurumisawa, J., Nakao, K., Billinghurst, M., Airola, C., Kato, H., et al. (2000). Augmented Groove: Collaborative Jamming in Augmented Reality. *SIGGRAPH'2000 CA&A,* pp. 77.

Rekimoto, J. (1988). Matrix: A Realtime Object Identification and Registration Method for Augmented Reality. *In Proceedings of APCHI'98*. ACM

Rekimoto, J., Ayatsuka, Y., Hayashi, K. (1998). Augment-able reality: Situated communication through physical and digital spaces. *In Proc. ISWC'98*. IEEE.

Rekimoto, J., Nagao, K. (1995). The World through the Computer: Computer Augmented Interaction with Real World Environments. *In Proceedings of the UIST'95*, ACM, pp. 29-36

Rekimoto, J., Saitoh, M. (1999). Augmented surfaces: A spatially continuous work space for hybrid computing environments. *In Proc. CHI'99*, ACM, pp. 378-385.

Schmalsteig, D., Fuhrmann, A., Szalavari, Z., Gervautz, M. (1996). Studierstube - An Environment for Collaboration in Augmented Reality. *CVE '96 Workshop*

Schmalstieg, D., Fuhrmann, A., Hesina, G. (2000). Bridging multiple user interface dimensions with augmented reality systems. *ISAR'2000*, IEEE, pp. 20-29.

Ullmer, B., Ishii, H. (1997). The metaDesk: Models and Prototypes for Tangible User Interfaces. *In Proceedings of the UIST'97*, ACM, pp. 223-232.

Ullmer, B., Ishii, H., Glas, D. (1998). mediaBlocks: Physical containers, transports and controls for online media. *SIGGRAPH'98*, ACM, pp. 379-386.

Wellner, P. (1993). Interaction with paper on the digital desk. *Communications of the ACM,* 36(7), pp. 87-96.