The Interaction Table – a New Input Device Designed for Interaction in Immersive Large Display Environments

M. Hachet and P. Guitton

Labri, Universite Bordeaux1, Bordeaux, France

Abstract

Large display systems such as Reality Centers or Powerwalls, allow several users to be immersed in a virtual environment while being located in the same physical space. The characteristics of such systems induce new problems and new constraints as far as it concerns the interaction. According to the lack of input devices well adapted to large displays, we are developing a new interactor: The Interaction Table. This device, composed of a movable tray fixed on a pillar, offers 6 DOFs and uses both isotonic and isometric information. The table top offers a 2D plane on which the position of a pen can be recovered. Many 2D and 3D interaction techniques can be used to accomplish the different interaction tasks (navigation, manipulation, selection, system control) dealing with different space ranges. The design of the Interaction Table makes it accurate and easy to use without any effort. Its auto-supported aspect makes it a non constraining tool, which can be shared by all co-located users. We illustrate the utility of the Interaction Table through a real application of 3D geomarketing.

Categories and Subject Descriptors (according to ACM CCS):
I.3.1 [Computer Graphics]: Input devices
I.3.6 [Computer Graphics]: Interaction techniques
I.3.7 [Computer Graphics]: Virtual reality

1. Introduction

The necessity to work together around specific projects has led to the creation of collective equipments for virtual reality (VR). On one hand, computers equipped with individual displays, such as classical monitor desktops or HMD systems, have been interconnected to offer a distributed or distance collaboration between several users: this corresponds to Collaborative Virtual Environments (CVE)[5]. On the other hand, very large displays (several m2) have been developed to allow several people located in the same physical space to be immersed in a unique virtual environment (VE) (Figure 1). The users can collaborate face-to-face, this makes the large displays a very useful and effective tool for collaborative tasks, particularly when decisions have to be made. For example, in the automobile industry, different specialists (e.g. engineers, designers, ergonomists) can collaborate face-to-face in front of a large projection screen in order to design a new car by interacting with a virtual model. The use of large displays in automotive design has been studied by Buxton et al. [3]

Figure 1: *Hemicyclia, an example of large display.*

Despite the fact that the use of large displays increases, well designed input devices for interaction with such systems do not exist. Usually, different people are immersed in the same virtual environment; however, only one user sitting behind a console interacts with a mouse. The majority of input devices adapted to immersive VEs has been developed according to HMD systems (e.g. gloves, tablets and pens). These devices can sometimes be used with systems

using other display devices for visualization, but very few approaches try to directly increase the performance of users collaborating face-to-face in front of large displays. Some authors have been interested in interacting with large displays used as whiteboards. For example, in [6], a pen is used to perform classical 2D tasks like selecting a menu or moving an item. Currently, many projects are based on laser pointers [12] to accomplish these 2D tasks, but as far as we know, it does not exist any input devices specially designed to accomplish 3D interactive tasks in a large display context.

The lack of input devices well adapted to such systems motivated us to develop a new device: the Interaction Table. The contribution of the Interaction Table is to allow several users to effectively interact with VEs displayed from projections on large screens. Many 2D and 3D interaction techniques can be used to accomplish the different interaction tasks (navigation, manipulation, selection, system control) dealing with different space ranges. The Interaction Table is a non constraining and effortless mechanical system offering a passive haptic feedback. Due to the appropriate correspondance between the mechanical affordances and input modes, a long learning period is not required.

We have developed a first prototype using existing commercial components, including a joystick and a spacemouse, in order to validate the main concepts. This paper describes a list of fonctionnalities, which will be performed in the version we are now developping.

In the first part of the paper, we recall the characteristics of large display systems. We demonstrate that none of the classical existing input devices are well designed for such systems. In the second part, we describe the Interaction Table and show its advantages. In the last part, we provide an example of the use of the Interaction Table in a real application of 3D geomarketing.

2. Large Displays

Large display systems can be used as an alternative to HMDs to obtain immersive VEs as it is shown in [14]. Among these systems we can cite Reality Centers, Vision Domes, Powerwalls and other systems using large surfaces of projection or backprojection. Developing an interactor adapted to large displays implies a prior study of the characteristics of such systems. The main characteristics of large displays are:

- Large size of the displays
- Collective experiment
- Unconstrained movements
- Visualization of the real environment
- Local and global information in the same image

Large displays offer good visualization for several users (between 10 and 20 at the same time), who can collaborate in the most natural way due to their physical proximity. The

users can move in a large area and can see their own bodies as well as the bodies of the others. The users can visualize the general virtual scene or focus on a part of it. According to these characteristics, we will now define the ideal properties an interactor for large displays should have.

2.1. Large size of the displays

Actual technologies allow to display one or a few views of a virtual world at the same time onto a large surface of projection. However, many users can visualize the same virtual world when facing a large display. Each user has his/her own point of view, which is different from the camera point of view (Figure 2).
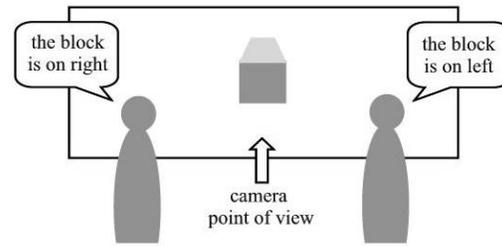

Figure 2: *Large displays induce different points of view.*

The camera point of view can be attached to a user by means of tracked systems. However, as the goal of large display systems is to offer a good visualization to a majority of users, the camera point of view is often fixed onto the normal of the screen. Though, several users can be immersed in a VE while keeping a quite satisfying view on it.

When using classical monitor desktops or HMD systems, both the user and the camera points of view are the same. The user can interact with the VE by using many interaction techniques. When using large display systems, one has to ask himself if this interaction techniques can still be easily used according to the fact that the position/orientation of the user can be different from the position/orientation of the camera point of view. Many interaction techniques are totaly dependent on the point of view (eg. selection by ray casting towards the object with a virtual pointer). Others, are not dependent on the point of view (eg. selection by pointing the projection on the screen of the object with a laser pointer). Our objective is to propose an input device from which a huge number of interaction techniques can be used. Therefore, the techniques, whether dependent or non dependent on the point of view must be supported.

2.2. Collective experiment

At present, a master of ceremonies usually interacts with the VE (usually with a simple mouse) while others ask him to

accomplish different tasks. We strive to make possible the interaction between the VE and all users facing the screen. Two options are possible: equipping all users susceptible to interact with a device or having one device that can be used by everyone.

The first solution comes in the scope of SDG (Single Display GroupWare), where the problem is to allow the interaction between the computer and each of the co-located users equipped with their own input devices. This strategy induces problems of conflict and interferences. Moreover, according to the discussion above, the interaction techniques that will be able to be used must not be dependent on the point of view. Some works are done with laser pointers, PDAs and speech and gesture recognition [22]. Such devices do not offer interaction techniques well adapted to 3D interaction but they constitute a potential as it concerns the collaborative tasks.

The second solution makes the use of devices more adapted to 3D interaction possible, but it can be quickly noticed that it is not easy to give each other these devices. This is evident for glove systems. Concerning the free moving devices like a Wand or other tracked systems, giving each others the interactor may be more convenient than with a glove, but it is certainly not completed easily. Furthermore, large display showrooms are often dark. The most reasonable way seems to be the use of an auto-supported device (i.e. nobody has to carry it), that avoids the device acquisition problem. According to the discussion "large size of the displays", the position of this device should be fixed to the position of the camera point of view in order to offer the user being interacting with the VE the best view on the virtual world.

2.3. Unconstrained movements

A substantial advantage of large screens is that users can move in a large space without being restricted by any wires. By developing a new device for interaction, we have to be careful not to decrease this freedom. One of our objectives is to propose a non constraining device because the more free the users will be, the better they will feel. Users must be able to "enter" and "exit" the virtual environment without being limited by any materials. Even if interaction is an essential component of virtual reality, it is not a goal. The interaction tasks represent a small part of the total time of a regular use of VR applications. The other tasks consist on visualization, understanding of the data, decision making... Users do not need to carry anything when they are not interacting. It is clear that an input device, particularly in a large display context, must not be constraining.

2.4. Visualization of the real environment

Contrary to HMDs, the use of large screens enables the view of the real environment, including people and equipments. This characteristic has to be taken into account concerning the development of a new device. For example, using virtual hands as an interaction metaphor may be confusing because the user seems to have four hands, as van de Pol et al. noticed in their study about interaction using a workbench [19]. Moreover, using a free space device can be problematic. Effectively, if the user has to move the device in front of his face, the display becomes partially occluded, and the immersive feeling is affected. This problem is not encountered when using an auto-supported device.

2.5. Local and Global Information in the same image

Systems based on large displays are the only ones that offer a local and global visualization in the same image. This suggests, that an ideal interactor should be effective for large movements and be accurate for small ones.

2.6. Further considerations

We also have to consider, that large displays can be used in many contexts including manipulation of objects in the personal space and navigation in the vista space. The vast majority of input devices are well suited to operate within the range of personal space [18]. This is the more efficient space for HMDs. Very few of them can be used in other virtual spaces. Furthermore, the majority of existing devices are effective for one task, but they can rarely be used for other tasks (for example, lasers are mainly selection tools). Moreover, large displays can be used with or without stereo viewing. Thus, we cannot develop a device which would be based on touching or grabbing directly the virtual objects. Our objective is to propose an input device that can be effective in a majority of contexts where persons are facing a large display. We are not aiming to find the best interactor adapted to one specific task. From this study, and from a review of existing devices, we propose a new input device adapted to large displays: the Interaction Table.

3. Hardware an software description

The Interaction Table looks like a table (Figure 3). The top is a circular tray which can pivot around the central pillar. This gives 2 DOF for the rotational information. The third rotational DOF is given by a wheel turning around the tray. The positional information is obtained by moving the tray in the desired direction, relative to its orientation. Therefore, the Interaction Table is a 6 DOF device mixing isotonic and isometric information: rotating the tray is performed by using the free moving component of the Interaction Table (without resistance or with a constant resistance). The rotations are performed through a position control process. Translating the tray occurs due to pressure which allows the tray not to move far away from its base. Translations are performed through a rate control process (Figure 4). The top of the table is equipped with a pad which allows capturing the position of a pen. Buttons situated on the top of the table can

be used to accomplish actions. The resistance of the tray can be controlled.
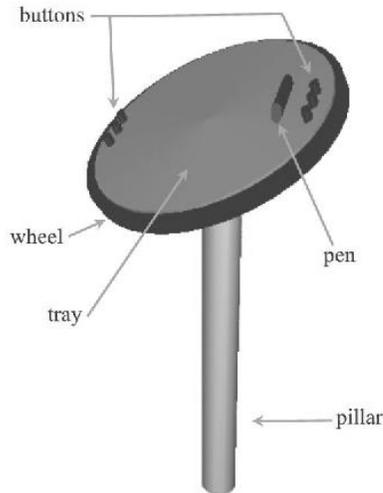

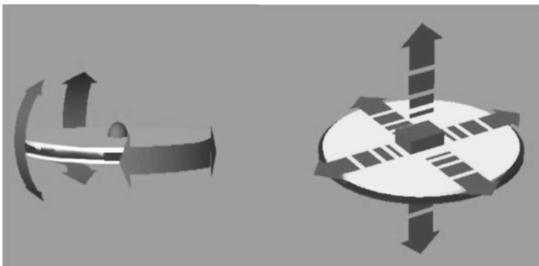**Figure 3:** *The Interaction Table.*

Figure 4: *Movements of the Interaction Table. The orientations are given from the isotonic component (on the left) while translations are given from the isometric component (on the right).*

We built a prototype composed of a Microsoft SideWinder joystick (first two DOFs), and a movable disk on which a mouse is fixed to recover the third rotational information. The transitional information is given by a Logitech Magellan Space Mouse fixed on the movable tray. We experimented the recovering of 2D points on the table top by using an Aiptek HyperPen tablet. A system of cables controls the resistance of the joystick (Figure 5).

Our test environment is a SGI Reality Center, composed of an Onyx2 (6 R12000, 3 IR3 engines), 3 Barco video projectors (BarcoReality 909), and a hemicylindric screen (3m x 10m).

The software part is developped by using WorldToolKit.

Figure 5: *Our prototype.*

It consists of the implementation of some (few) interaction techniques in which a semitransparency representation of the table top can be either displayed or absent. The similar shapes of both the virtual and the real table top make the user feel as if he were holding the virtual table. The advantages of the use of semitransparency in 3D human interaction have been highlighted by Zhai et al. in [24].

4. Advantages of the Interaction Table

If we consider the discussion above, it is clear that the Interaction Table is adapted to large displays. Everybody can interact with the VE, and then move in the area without being constrained by wires. The Interaction Table appears to be a control centre accessible by everybody. Similar to a helm permitting to drive a ship, the Interaction Table is a reference point in the room from which it is possible to interact with the virtual environment. Its fixed position never occludes the screen.

4.1. Interaction techniques

The Interaction Table permits the use of different interaction techniques to accomplish different tasks. We present here these techniques.

4.1.1. Navigation

The navigation, or more precisely the travel component of navigation, can be performed by using the Interaction Table as a 3D wheel. According to Ware and Osborne [20], this corresponds to a Flying Vehicle Control metaphor. The user

moves the table top with his two hands and he sees what he would see if he were sitting on the table. The user can give an orientation to the plane he is holding and then apply translations on it, or he can apply both kinds of transformations at the same time. The orientation of the table top can be fixed in order to allow only translations in one plane (for example, using the horizontal plane for a navigation in a city). This corresponds to a walking metaphor. Adding constraints can increase the performance in Immersive Virtual Environments [2].

Another interaction technique for navigation consists on selecting the endpoint of the travel trajectory. This can be made for example by selecting a 3D point (see Selection paragraph) or by choosing a location out of a list (see System Control paragraph). Moreover, having a 2D plane offers the possibility to use 2D interaction techniques for navigation. For example, in [1], vectors are drawn on the 3D Palette to indicate the direction to follow.

Ware and Osborne proposed to change the point of view by moving the whole scene. This is the Scene in Hand metaphor. The Interaction Table is particularly well adapted to this interaction technique.

4.1.2. Manipulation

The Interaction Table has a substantial advantage by offering a physical support. If we consider the Scene in Hand metaphor, the user holds the scene through the tray. The orientation of the scene is the same as the orientation of the tray. This provides an important feedback. We implemented a method which consists of, first, moving a representation of the table top into the scene, and then, after pressing a button, attaching the scene to the virtual tray. This allows fixing the point of rotation. In a huge scene, we can first be interested in the whole area, then in details. Moving the virtual tray permits one to move his centre of interest.

As it concerns the manipulation of single objects, the objects must first be attached to the table top movements (by a selection process) and then, the user holds the object in his hands (Figure 6). The benefits of using physical props has been highlighted by Hinckley et al. [10]. After Hoffman [11], touching virtual objects physically using tactile augmentation enhances the realism of VEs. Ware and Rose showed, that the shape of the physical tool did not influence the performance [21].

Moving the tray according to its orientation causes a coherent stimuli-response. Assuming that users are visualizing a virtual car attached to the movements of the Interaction Table, they can quickly change the orientation of the car in order to visualize it without clutch systems and without twisting one's arm. If they want the car to be translated in its local frame, they just have to push the tray in the desired direction, according to the cubic mouse [8].

Manipulation of objects can be performed by using the
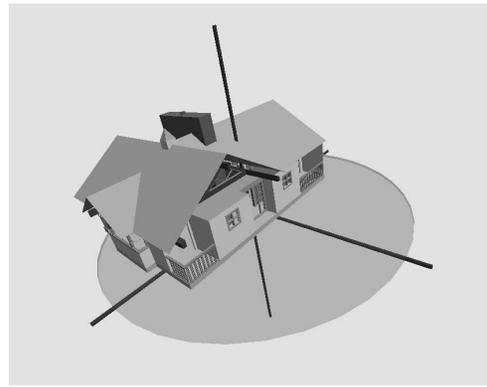

Figure 6: *The virtual tray is attached to an object.*

2D plane, too. For example, after having positioned the representation of the table top in the 3D space, objects can very accurately be translated by a slide of the pen on the tray. The user can perform translations in a 2D plane which is easier than in a 3D space.

4.1.3. Selection

By displaying a representation of the tray, we can use different methods to select an object (or a group of objects) with the Interaction Table. The first of them consists on using the virtual tray as a 3D pointer. The size of the virtual tray can be adapted in order to modify the accuracy.

Secondly, the virtual tray can be used as a butterfly net which corresponds to 3D sweeping as it is done in [16].

Another method is to select 2D points on the tray with the pen. The virtual tray is positioned in the 3D space and has exactly the same orientation as the physical one. Therefore, it is intuitive to select a point on the virtual tray corresponding to the location of the pen on the physical tray. It is important to notice that the selected 3D point may correspond to none of the virtual objects. A 2D description language can be used to select an object or a group of objects which are located on the virtual tray. This can be performed by circling an area for instance. This last method can be applied to applications which are based on a horizontal plane such as GIS. For such applications, this interaction technique can be very effective.

A last approach is the use of the virtual tray as a 2D plane through which the user sees the virtual world. The physical tray in front of the user represents the screen in miniature. Therefore, it is possible to directly "touch" the screen through the table top. From this way, it is possible to choose objects by selecting their projection on the virtual tray.

4.1.4. System Control

System control (e.g. use of menus, typing values and texts) is often hard to perform in a 3D virtual environment. Input

devices which are most adapted to these tasks are tablet-and-pen systems [15] because they offer 2D physical support. All techniques used by tablet-and-pen like devices can be used with the Interaction Table. Writing annotations or choosing colours out of a list is no longer difficult. The techniques used with classical monitor desktops can be used with the virtual tray, the table top replacing the desk. The user positions his "virtual monitor" anywhere in the 3D space.

It can be noticed that having a 2D plane in hands permits to sketch virtual objects by using revolution, extrusion or other methods. This is illustrated in [7] through the translucent sketchpad.

We have seen that the Interaction Table allows a large set of interaction techniques. It can be used as a support for other approaches. For example, a WIM [17] can be attached to the table top, and the interaction techniques proposed by Stoakley et al. can be adapted to the Interaction Table.

4.2. Design

The design of the Interaction Table offers some substantial advantages.

4.2.1. Effortless

Contrary to free space input devices, users do not have to carry anything. Carrying a device is tyring as it has been noted in [17]. Further than carrying a physical object, we think that moving one's hands in the space, as it is necessary with gesture recognition, may become fatiguing and annoying. The Interaction Table offers a support giving the user a rest. The auto-supporting aspect of the Interaction Table works without encountering problems caused by the device acquisition.

4.2.2. Mechanical system

In its actual version and the versions we plan, the Interaction Table is a mechanical system. A mechanical system has the advantage of being accurate, responding with a short lag, not producing interferences, and being inexpensive. MacKenzie and Ware [13] show the importance of the lag in interactive systems. The system of cables we use gives the Interaction Table a physical memory. This allows the user to interact with a VE, then do other things and return to interact with the VE without a change of the physical state of the Interaction Table. The fact that the Interaction Table has not moved induces that no modification of the virtual environment has been performed. The mechanical aspect of the Interaction Table allows one to view a force feedback system. In regard to the position control (rotations) and the rate control (translations), Zhai discusses the advantages of each of these approaches [23]. Among other things, he shows that position control allows one to quickly perform tasks whereas rate control is more accurate.

5. Interacting with an application of 3D geomarketing

We have developed a 3D geomarketing application [9] for a company: Cartegie [4]. We are going to see how the Interaction Table could be used in this context. Geomarketing aims to show strategical (e.g economical, demographical) information on geographical support in order to accelerate decision making processes (Figure 7).
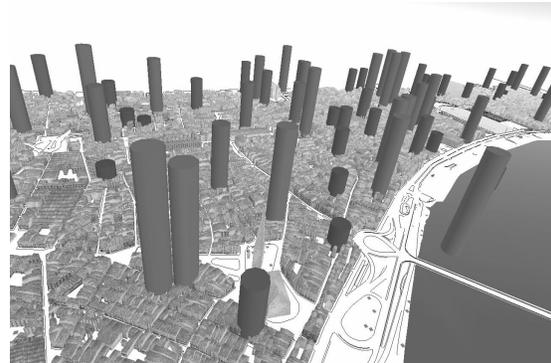

Figure 7: *Example of 3D geomarketing space.*

Using 3D scenes on large screens with stereoscopic vision helps to understand the data. However, as long as the interaction with VEs will be a problem, geomarketing will go without virtual reality. Effectively, tasks for geomarketing are various, and decision makers cannot waste time in performing these tasks. A very simple geomarketing scenario is presented in the following paragraphs.

At the beginning, we have a huge territory, but we only want to work on a restricted area. By pointing a precise location on the tray with the pen, we can travel to the corresponding point on the virtual disc. It is important to notice that this point can belong to an empty zone (i.e. we don't need to select a point which belongs to an object). Once arrived at the desired location by a travel process, we want to navigate in order to discover the zone. This can be done by using the Interaction Table as a 3D wheel. 1- We can use all DOFs for navigation, keeping in mind that the orientation of the tray is the same as the orientation of the viewpoint. This means we have a permanent mark according to the horizontal plane. 2- We can limit some DOFs. For example, by physically locking the orientation of the tray, we can travel only on the horizontal plane. This can be effective when navigating into an urban area.

Now, we want to visualize strategic data into the selected area. The first step consists on selecting the data we want to see. The Interaction Table can be used as a tablet in order to have a 2D plane allowing to select the data (Figure 8). For example, we choose to visualize the numbers of inhabitants per building by matching them to red cylinders. To comprehensively understand these data and their relation in the 3D environment, we have to watch them from different

Figure 8: *Use of the Interaction Table as a support to select items in a menu.*

viewpoints. This can be performed effectively and quickly with the Interaction Table as it has been described in the manipulation paragraph. From this perspective, we will be able to more accurately appreciate the differences between the heights of the different cylinders.

We decide now to position a new building in this area. This building can be sketched by means of the 2D physical plane. It can be easily positioned by sliding the pen on the table top.

Supposing we want more information concerning a particular building, we can select it by using the pen on the tray (by a pointing or a circling process). Alphanumeric information can be presented as a text written on the virtual disc, which can easily be located to offer a good reading. At the end, we want to annotate this building. This can be done by writing the text on the table top and positioning it where we want.

In this example, many tasks have been performed at different ranges of space. Throughout the sequence, any of the users was able to interact.

6. Conclusion

The Interaction Table is developed to fill the gap of input devices adapted to interaction with VE displayed from projection on large screens. Its design allows a variety of interaction techniques to perform several interaction tasks in different ranges of space.

The Interaction Table is not aiming to replace all existing devices. It constitutes a federative tool in front of a large screen for general nature classical interaction tasks. The specific interaction tasks can be performed by dedicated devices.

We have developed a first prototype of the Interaction Table using existing commercial components in order to

validate the main concepts. We are now developing, with mechanical and eletronic specialists, a second version for a comparative testing with existing devices. This second version will enable the definition of a uniform set of interaction techniques to be used with the Interaction Table.

The appropriate correspondances between the mechanical affordances and the input modes make the Interaction Table intuitive and easy to use. The auto-supported aspect allows the users not to be linked to the Interaction Table. This provide a feeling of freedom. Not feeling free in a VE is a common reason why people consider virtual reality as a complex tool that is closer to science fiction than to their real preoccupations.

Acknowledgements

References

1. M. Billinghurst, S. Baldis, L. Matheson, and M. Philips. 3d palette: A virtual reality content creation tool. In *Proceedings of VRST '97*, pages 155–156, October 1997.

2. D.A. Bowman and L.F. Hodges. Formalizing the design, evaluation, and application of interaction techniques for immersive virtual environments. *Journal of Visual Languages and Computing*, 10(1):37–53, February 1999.

3. W. Buxton, G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach. Large displays in automotive design. *IEEE Computer Graphics and Applications*, 20(4):68–75, July-August 2000.

4. Cartegie. *http://www.cartegie.com/*.

5. CVE. *http://www.crg.cs.nott.ac.uk/groups/cve*.

6. S. Elrod, R. Bruce, R. Gold, R. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, E. Pedersen, K. Pier, J. Tang, and B. Welch. Liveboard: a large interactive display supporting group meetings, presentations and remote collaboration. In *Proceedings of CHI 92*, pages 599–607, 1992.

7. L.M. Encarnação, O. Bimber, D. Schmalstieg, and S.D. Chandler. A translucent sketchpad for the virtual table exploring motion-based gesture recognition. In *Proceedings of Eurograhics '99*, pages 179–190, September 1999.

8. B. Frölich and J. Plate. The cubic mouse a new device for three-dimensional input. In *Proceedings of CHI 2000*, pages 526–53, April 2000.

9. M. Hachet and P. Guitton. From cadastres to urban

environments for 3d geomarketing. In *Proceedings of IEEE/ISPRS joint workshop on remote sensing and data fusion over urban areas*, pages 146–150, November 2001.

10. K. Hinckley, R. Pausch, J.C. Goble, and N.F. Kassel. Passive real-world interface props for neurosurgical visualization. In *Proceedings of CHI '94*, pages 452–458, April 1994.

11. H.G Hoffman. Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In *Procedings of VRAIS'98*, pages 59–63, March 1998.

12. D. R. Olsen Jr and T. Nielsen. Laser pointer interaction. In *Proceedings of CHI 2001*, pages 17–22, March-April 2001.

13. I.S. MacKenzie and C. Ware. Lag as a determinant of human performance in interactive systems. In *proceedings on Human factors in computing systems*, pages 488–493, April 1993.

14. E. Patrick, D. Cosgrove, A. Slavkovic, J. Rode, T. Verratti, and G. Chiselko. Using a large projection screen as an alternative to head–mounted displays for virtual environments. *CHI Letters*, 2:478–485, April 2000.

15. I. Poupyrev, T. Numada, and S. Weghorst. Virtual notepad: handwriting in immersive vr. In *Proceedings of IEEE VRAIS '98*, pages 126–132, March 1998.

16. D. Schmalstieg, L.M. Encarnação, and Z. Szalavári. Using transparent props for interaction with the virtual table. In *Proceeding of I3DG '99*, pages 147–153, April 1999.

17. R. Stoakley, M.J. Conway, and R. Pausch. Virtual reality on a wim: Interactive worlds in miniature. In *Conference proceedings on Human factors in computing systems*, pages 265–272, May 1995.

18. S. Subramanian and W.A. Ijsselsteijn. Survey and classification of spatial object manipulation techniques. In *Proceedings of OZCHI 2000*, December 2000.

19. R. van de Pol, w. Ribarsky, L. Hodges, and F. Post. Interaction in semi-immersive large display environments. In *Proceedings of EGVE '99*, pages 157–168, May-June 1999.

20. C. Ware and S. Osborne. Exploration and virtual camera control in virtual three dimensional environments. In *Proceedings of I3D '90*, pages 175–183, March 1990.

21. C. Ware and J. Rose. Rotating virtual objects with real handles. *ACM transactions on Computer–Human Interaction*, 6(2):162–180, June 1999.

22. M. Yoshida, Y.A. Tijerino, A. Shinji, and F. Kishino. A virtual space teleconferencing system that supports intuitive interaction for creative and cooperative work. In *Proceedings of I3D '95*, pages 115–248, April 1995.

23. S. Zhai. *Human Performance in Six Degree of freedom Input Control*. PhD thesis, University of Toronto, 1995. *http://vered.rose.utoronto.ca/people/shumin_dir/ papers/PhD_Thesis/top_page.html*.

24. S. Zhai, W. Buxton, and P. Milgram. The partial occlusion effect: Utilizing semi-transparency in 3d human computer interaction. *ACM Transactions on Computer–Human Interaction*, 3(3):254–284, 1996.