Chapter 7
Displays and Interaction for Virtual Travel

Anthony Steed and Doug A. Bowman

Abstract Virtual travel can be accomplished in many ways. In this chapter we review displays and interaction devices that can be utilized for virtual travel techniques. The types of display range from desktop to fully immersive and the types of interaction devices range from hand-held devices through to motion tracking systems. We give examples of different classes of device that are commonly used, as well as some more novel devices. We then give a general overview of travel tasks and explain how they can be realized through interaction devices.

7.1 Introduction

Being able to move the viewpoint in a virtual environment (VE) is a critical facility: small movements allow the user to get new perceptual cues to understand 3D space; larger movements allow the user to access different parts of the VE so that the user can experience them at a closer range, or access them if they were previously not accessible. Thus, when we consider VE systems, we can identify that travel has a range of purposes, which might require a range of different devices and interaction techniques. If one were building a training operator for a user of a control console, the user might be seated, and would only need to move her head to see different parts of the console, or look in different directions. If one were building a training simulator for fire evacuation, it would be important to be able to simulate movement through the large-scale environment.

A. Steed (✉)
Department of Computer Science, University College London,
Gower St, London WC1E 6BT, UK
e-mail: A.steed@ucl.ac.uk

D. A. Bowman
Center for Human Computer Interaction, Virginia Tech, Blacksburg, USA

F. Steinicke et al. (eds.), *Human Walking in Virtual Environments*,
DOI: 10.1007/978-1-4419-8432-6_7,

In what has been called the “virtual reality model” of interaction [30, Chap. 20] or the highest level of “interaction fidelity” [20], the user would just carry out actions in the VE in the same manner as she would in the real world. The system would track the body of the user and recreate virtual images, sounds and other cues that mimicked cues and sensations that the user would get from the analogous real situation. The user would see the VE from a first-person point of view and would be able to effect natural interactions with her body. If the user wanted to pick up an object, she could reach out her hand and grasp the object and lift it. More importantly for this book, if she wanted to pick up an object that was out of reach, she could walk over and pick it up.

Of course, the virtual reality model is an ideal because the hardware and software we use can’t simulate cues anywhere near as rich as the real world. Even if the participant isn’t walking or moving, our technologies for simulating tactile and force cues are very limited in that they can provide a few points of contact or small area of stimulation, whereas the task could involve the whole body of the user. If the user walks or otherwise moves, then there is a much more pressing problem: virtual reality devices only allow actual movement within a small area. This is for various reasons: the displays might be static (e.g., small room), tethered (e.g., a wired head-mounted display) or otherwise limited through infrastructure (e.g., a tracking system that only functions within a bounded region). We will discuss such technologies in greater detail later in the chapter, but there is a more pertinent question: how can we simulate real walking, giving the impression of unconstrained motion, when physical motion is actually limited? What devices can give the impression of walking on different surface types, over long distances or on different inclines?

There are two fundamental problems: walking is implicitly a task that involves very complex simulation of the walking surface, and walking involves inducing momentum into the moving object: the walker’s body. The first requirement might only be solved by what Sutherland called the *ultimate display* [33]. In his seminal paper he described that in this display the existence of matter would be controlled and that a speeding bullet could potentially be fatal. Thus the ultimate walking display would be a display that could simulate any surface by creating that surface. However even the ultimate display doesn’t directly solve the second problem: creating momentum in an object. Researchers are only just starting to solve the problem of configurable surfaces (e.g., see Chaps. 9 & 17 in this volume), and the problem of momentum is recognized and some attempts have been made to simulate it by pushing on the body (e.g., see Chap. 6 in this volume).

Reproducing natural walking is thus one of the toughest challenges in human–computer interaction. We can try to imitate real walking, but we will be limited in the range we can support, or the naturalness of the interaction. The alternative is to provide interaction techniques that produce movement, or travel, through the VE using other metaphors and devices.

In this chapter we outline the broad range of displays and devices that are used for travel techniques in VEs. Other parts of the book focus on reproduction of natural walking through sophisticated devices. We will place these in context of supporting the general task in a broad range of VE systems.

Section 7.2 will cover display systems for VEs.

Section 7.3 will cover tracking and interaction devices.

Section 7.4 will describe the range of travel metaphors that are commonly used.

7.2 Display Systems

Display systems for VEs come in many sizes, form factors and capabilities. We distinguish a few types of display:

- Desktop displays
- Wall-sized displays
- Surround-screen displays
- Head-mounted displays
- Mobile augmented reality displays
- Hybrid situated displays

By a desktop display we mean a display such as a humble monitor or TV of the type that is found in most offices or homes. It may show a VE, or other media, but it does not cover a large proportion of the user's vision compared to other displays. The display might use stereo and/or head tracking, but the display doesn't surround the user, so it can only provide a small window into the VE. Obviously one can represent self-motion on such a display. Indeed, many modern video games focus on travel and exploration around a VE, but many of the perceptual cues received are unlike the ones received from actual self-motion. For example, the cues do not extend into peripheral vision.

Once a display becomes very large, such as a single wall of a room (Fig. 7.1), then it covers much more of the user's vision. One characteristic of such a display is that nearby items such as other humans can be depicted "life-size" (with realistic scale) and thus the representation of self-motion can exploit the fact that the user will be able to judge heights more accurately, since he can have an eye-level that can corresponds to his actual height. The displays might be flat (e.g., an actual room wall) or they might be large curved screens. A key characteristic is that the display system is not fully surrounding, so it is possible for the user to turn away from it. Thus any walking interface using such a display would normally require some facility to turn the viewpoint.

At this scale, it is much more common to provide stereoscopic viewing of the display. There are several technologies for this such as shutter glasses, polarising glasses, or color filters. Each works by presenting separate images to the left and right eyes. Most displays of this type support only a single user with head-tracked stereo, but multiple users can be supported [1]. The current state of the art is the C1X6 (Fig. 7.2), a prototype display that supports six users, each with a stereo view from her own viewpoint [15].

The next type of display is a surround-screen display (SSD), of which the most common type comprises multiple large flat display surfaces. These displays are often

Fig. 7.1 An example of a wall-sized display. The EVEREST display is 30 ft long by 8 ft tall and displays 35 million pixels. Image courtesy of the National Center for Computational Sciences, Oak Ridge National Laboratory

Fig. 7.2 The C1X6 supports six users, each seeing a first person stereo view [15]. Courtesy of Virtual Reality Systems Group, Bauhaus-Universität Weimar, Germany

Fig. 7.3 *Left* UCL's four-walled CAVE™-like display. *Right* looking down into the virtual pit

referred to as CAVE™-like displays (Fig. 7.3, left), named after the original presentation of such a display [6], and the current trademarked commercial version from Mechdyne Corporation. Again it is hard to draw a line between large wall displays and SSDs as the most common type of SSD is a CAVE™-like display with four sides (front, left, right, and floor). Such displays do not completely surround the user, but when standing in the center of such a system and facing forward, the user's vision is almost completely filled with the rendered display of the VE. Importantly, in the most common four-screen configuration the user has peripheral vision of the VE to his left and right, and also down. Thus the visual cues for motion are more powerful. The "virtual pit" scene that is often used as a demonstration of the effectiveness of virtual reality [38] is especially powerful in a CAVE™-like display with a floor because the drop is shown near the feet of the user (Fig. 7.3, right).

Six-sided CAVE™-like systems (Fig. 7.4) do exist, though there are only a handful in the world. The engineering of a fully surrounding system of this type is complex because the floor must be a back-projection screen while being safe for groups of users to stand upon.

Many variations of SSDs exist in all sorts of display configuration. Domed displays are common in certain forms of entertainment (e.g., planetariums), though these are typically used for audiences in the dozens rather than individuals. More exotic displays include the Allosphere (Fig. 7.5) at the University of California, Santa Barbara, a spherical display that supports moderately sized groups [12], and the Cybersphere (Fig. 7.6) [9], a spherical display that supports one person and also acts as a treadmill.

Head-mounted displays (HMDs) are the type of display that is most commonly associated with the phrase "virtual reality." While the basic technology has been available for over thirty years, recently there has been a resurgence in interest and new displays have come to the market. The advantage of HMDs is that they can create the impression that the VE fully surrounds the user: as he turns his head the visual and audio displays can be updated to reflect that motion. This requires some form of head tracking, whether built into the HMD or attached to it. There are many varieties

Fig. 7.4 HyPi-6 display at the Fraunhofer IAO, Stuttgart, Germany. Courtesy of Oliver Stefani

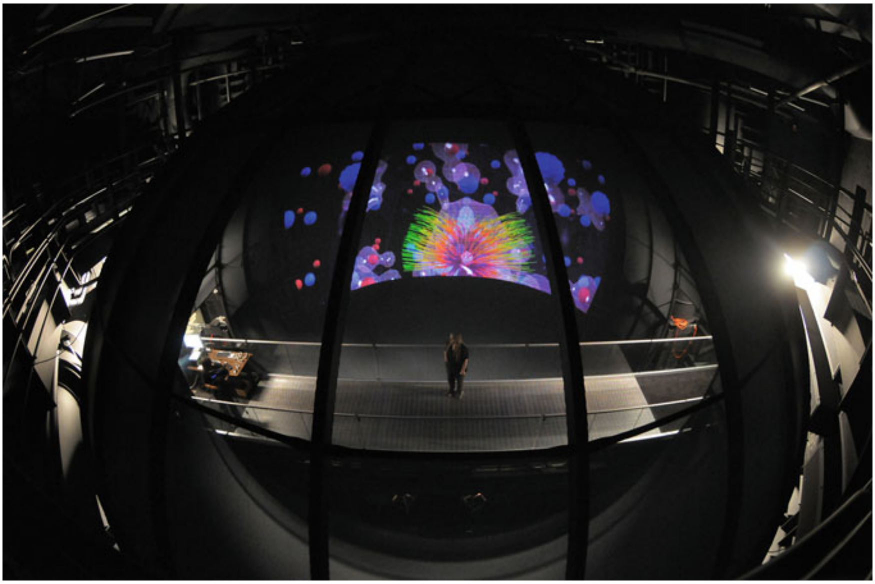

Fig. 7.5 Director JoAnn Kuchera-Morin on the bridge of the AlloSphere, photo taken from above the instrument. Photography by Paul Wellman. Courtesy of University of California, Santa Barbara

Fig. 7.6 Cybersphere. Courtesy of Professor Vinesh Raja, Director of Informatics and Virtual Reality, International Digital Lab, WMG, University of Warwick

Fig. 7.7 *Left* Sony HMZ-T1. *Right* nVis SX111. Courtesy NVIS Inc

of HMD: not all provide separate images for the left and right eyes; they vary in resolution; and, most importantly, they vary in field of view (FOV). There are many older HMDs with very low resolutions and FOVs, but a modern HMD is the Sony HMZ-T1 (Fig. 7.7, left) released in late 2011, which used OLED panels, had twin screens at 1280 × 720 color pixels and a field of view of 45°. At the more specialist end of the market, the nVis SX111 (Fig. 7.7, right) has LCOS panels, 1280 × 1024 color pixels per eye, and a field of view of 111°.

A useful feature of HMDs is that since they are mounted on the user, the user can walk about. Thus a HMD can be placed in a relatively large space given a tracker that can track the user over that space. In practice this is limited by cabling, but wireless HMDs are now available.

Augmented reality (AR) displays are used to present virtual elements within a view of a real environment. That is, the VE is seen as an augmentation of or intervention into the real environment and both are to be understood by the user as a single consistent environment. We will focus on two sub-types: *mobile augmented reality* and *hybrid situated displays*. The key features in both are that the user walks about the "display" (where the real world is considered to be part of the display), and the display is typically larger than that provided by other technologies we have discussed. In particular, these two types of technology allow walking around much larger spaces.

Mobile augmented reality (Fig. 7.8) typically uses a see-through HMD or a hand-held display, which allows the user to see the real world with graphics overlaid on it. Early demonstrations of these systems included annotations that appeared fixed in place as the user walked [8]. In contrast, the *mixed-reality* project was an example of a project that attempted to create the impression that the physical environment had been extended with new buildings [36].

A hybrid situated display consists of a physical environment with a variety of situated displays within it. A good example is the Infantry Immersion Trainer (Fig. 7.9), a military training facility where soldiers walk around a warehouse space that has been converted into a physical mock-up of a Southwest Asian town [25]. Within this space, for example in mock house interiors, are large screen displays that show life-sized characters that the soldier must react to. The facility thus acts as a type of advanced shooting range. For our purposes the key feature is that the display, being the combination of physical props and VEs, comprises a fully immersive environment around which the participant walks. Theme park experiences are often of this nature as well.

Some observations on the affordances of these display types for walking in VEs are presented in Table 7.1.

7.3 Interaction Devices

In supporting walking interfaces, we can identify two important tasks:

1. Orientation of the viewpoint in the VE
2. Movement of the viewpoint in the VE

With real walking both are specified by the muscles and state of the body including legs, spine bend, neck orientation and eye rotation. As users move, each of these changes constantly in a complex manner, but we note that orientation of the viewpoints (i.e., the eyes) can be controlled separately from the gross direction of movement. We also note that while walking the eyes are in constant motion, not

Fig. 7.8 A later version of the touring machine system. Courtesy of Steve Feiner, Columbia University

just in the direction of travel, but up and down. All of these, and many other subtle effects, need to be recreated in VEs if a purely natural walking interface is desired.

In the virtual reality model discussed previously, the concept is that the system tracks the body and recreates the perceptual cues in a structure that is analogous to the real world. In practice this means that the system tracks the head of the user and uses this to orient a pair of virtual cameras, one for the left eye and one for the right eye. It would be ideal to track the eyes as well, both direction and focus, and then adapt to that, but that is beyond the current state of the art.

Since it is not possible to track the user in an unlimited area, interaction devices must be used to provide control input to effect the two tasks we identified above. While we will discuss the actual interaction tasks in the next section, it is worth pointing out the degrees of freedom that are required. The task of placing a camera in the world involves six degrees of freedom (DoF): three for orientation and three for translation. For two eyes, two cameras are required, but typically these are rigidly attached to the head, so a single six-DoF position/orientation needs to be calculated. In one extreme, say the desktop display type from Table 7.1, nothing about the user's

Fig. 7.9 Marines from 3rd Battalion, 1st Marines, confront avatars, or virtual humans, while clearing a room at the Office of Naval Research Infantry Immersion Trainer. U.S. Navy photo by John F. Williams from WikiMedia Commons

Table 7.1 Comparison of display technologies

Type	Viewing type	User own motion	Notes
Desktop	Through the window/Low FOV	None (typically) or very limited (<1 m)	Usually no head tracking, user typically seated
Wall-sized	Through the window/Medium FOV	None through to limited (<3 m)	Commonly a single person can be tracked. For a group, no tracking used
Surround-screen displays	Partially–fully surrounding/High FOV	Limited (<3 m) to moderately large (∼10 m)	Commonly a single person can be tracked. For a group, no tracking used
Head-mounted displays	Fully surrounding/ low-high FOV	None through to wide area	Single user only
Mobile augmented reality	Fully surrounding/ FOV of virtual world limits	None through to wide area	Tracking quality and thus registration vary enormously
Hybrid situated displays	Fully surrounding/ FOV not applicable	None through to wide area	Tracking might be localized to certain areas of the display

actual head position is known, so the system must provide a metaphor for controlling the viewpoint using devices such as joysticks, mice, keyboards, etc. At the other extreme, such as mixed-reality displays, because the user can move through the full extent of the VE, head tracking is sufficient to give us the relevant viewpoint. In between there might be a hybrid: tracking the head over its limited range of motion, and providing a separate control mechanism for movement over longer distance.

We can thus identify several types of interaction device that might be used:

- Pose & movement sensors
- Position trackers
- Hand-operated devices
- Hybrid devices

As with the display types identified in the previous section, the distinctions between these types are quite subtle. Furthermore, any one device might incorporate elements of two or more of these types (e.g., a handheld device that contains a joystick and a position tracker). We also note that there are hundreds, perhaps thousands, of devices that are available for interaction in VEs. We'll thus pick some key examples that are either historically important or are very widely available and thus commonly used.

We start with pose & movement sensors (also called inertial sensors) because they are commonly available and cheap. This type of sensor is usually an integrated unit that can sense orientation and some types of movement: three DoF for an accelerometer that detects linear acceleration (including gravity, which can be used to estimate pitch and roll orientation), three DoF for a gyroscope that measures changes in rotation, and three DoF for a compass that can give an absolute heading. In 2012 a modern smartphone would typically include all of these sensors and a GPS unit to give ~5–10 m accurate global position information. This data is not itself sufficient to generate, for example, a precisely registered augmented reality (Fig. 7.10), as the sensors are not very high quality and do not give precise readings that are registered in local coordinates. It is not generally possible to build a local position tracking system from an accelerometer because the sensors drift over time. In addition the

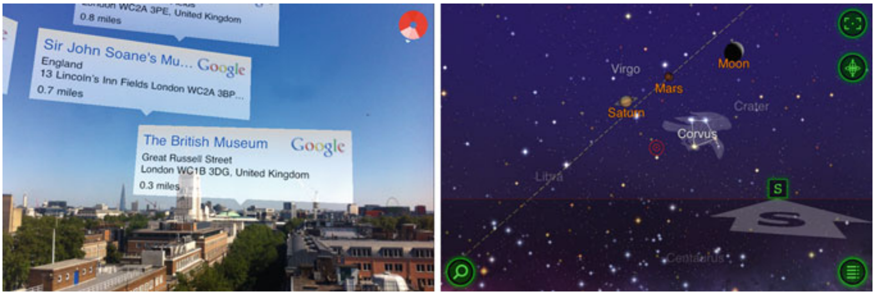

Fig. 7.10 *Left* acrossair browser, acrossair, showing museums south from University College London. *Right* Star Walk, Vito Technology, Inc

Fig. 7.11 The eMagin z800 helmet contains a three axis rotation sensor

accelerometers cannot tell the system the height above the ground or distance from any local feature without other calibration. Calibrating against GPS fixes is problematic because those fixes are inaccurate and the civilian implementations include a moving offset by design. Thus currently, outside of the lab, mobile augmented reality is limited to display modes that use gross position referencing, such as labels over whole buildings, rather than being able to accurate align an object to something the size of a window or door. Researchers that are prepared to integrate more equipment than can be found on current smartphones can build systems that maintain good registration, but the vision of being able to walk around and see an augmented real environment where the augmentations are seamlessly integrated is still a few years off.

The other common use for pose & movement sensors is for head-mounted displays and handheld devices. Some consumer HMDs include a two-DoF or three-DoF rotation tracker so that the user can rotate her head to directly control the orientation of the viewpoint (Fig. 7.11). These sensors do not track linear acceleration, so they cannot provide motion parallax, thus in the lab these sensors are usually integrated with a position tracking technology as well. Recent game controller devices such as the Nintendo Wii Remote and Sony Move controllers integrate accelerometers and gyroscopes. Again, these two controller technologies can only tell the console about local rotation of the device. This is insufficient to do position tracking of the technology, so it can't, on its own, support pointing at the screen. Both provide separate position tracking for this.

The basic component sensors thus form an important part of more sophisticated trackers. For example, the InterSense IS900, which is a current state of the art position tracker that is commonly used in virtual reality laboratories, includes an accelerometer as a component. Accelerometers in combination can create more sophisticated models of moving objects. The XSens system is a motion capture technology where the user wears several accelerometers. From the relative acceleration of the limbs a model of the skeleton of the user can be built. However, over time, the position of the user will drift if only the accelerometers are used to calculate his location.

Although we have already seen some crossover, we distinguish position trackers from the sensors already discussed because they return an absolute 3D position in a fixed coordinate system. That is, if the tracked device is moved and returned to the same position, the tracker reports the same position up to its own tolerance for precision and accuracy. This is unlike an accelerometer that will drift over time. Position trackers are a well-studied component of augmented reality and virtual reality. The technology changes relatively slowly, so previous surveys [21, 41, 44] give good overviews that are still very relevant. We thus detail the most common technologies in use and highlight their strengths and weaknesses.

For small spaces, the stalwart of the field has been the magnetic tracking technologies epitomized by the Polhemus Fasttrak. This can return six DoF of a sensor that is attached to or embedded in a control device. A standard unit might track 1–4 such sensors, and a common set up would be to attach one to a HMD and embed a second in a hand-held controller. Because the tracking is magnetic, there is no need for the sensor to have a line of sight to the base station, as with optical tracking (see below). Thus the trackers are commonly used in situations where occlusion is likely. However, the sensors typically do not work over spaces of more than 3 m by 3 m (the space is often less than this), and are affected by metal in the environment.

A cheaper, but more limited technology is the visual tracking systems that can track the head, hands, or other body parts in a small volume. In particular, the Wii Remote contains an IR sensor that can track a bar of LEDs and thus can estimate the relative position of the WiiMote from the bar. This is used to allow direct pointing at the screen. The Sony Move controller uses a camera near the display to track the 3D position of a large light source on the controller. Microsoft's Kinect uses a depth camera to track the skeletons of one or more people in front of the display. Alongside these three currently popular technologies, there are quite a few others that have the aim of giving a limited range of direct movement control in front of the display (e.g., NaturalPoint TRACKIR, Logitech Head Tracker). Most of these track only position, not orientation (although the Wii Remote and Sony Move controllers add pose and movement sensors to measure orientation indirectly), and all of them assume that the user is facing in the direction of the display.

In situations where larger tracking volumes, integrated position and orientation tracking, and greater flexibility of movement are needed, high-end optical tracking systems similar to those used in motion capture are frequently used (Fig. 7.12). Common systems include Vicon, OptiTrak, ARTrack and PhaseSpace. The first three systems use passive markers that are mounted on a device, an item of clothing or a full body suit. The passive markers are retro-reflective, and the camera has an infrared light source to illuminate the markers. The PhaseSpace system uses active markers similarly arranged. An important part of the technology is that the system tracks the positions of the individual points in 3D space, but does not itself track the orientation of the marker. Thus multiple markers on a rigid or near-rigid object are needed to determine the orientation of the object. Position and orientation tracking

Fig. 7.12 A user wearing a motion capture suit within an OptiTrak installation. The user is taking part an experiment involving a simulation of a train. Courtesy of Angus Antley, University College London

of rigid bodies can be quite accurate and stable, so this approach is used for head tracking in order to ensure stability of the view. In some situations it is useful to have a full motion-capture setup where the body of the user is captured to monitor her limbs movements and potentially create a virtual representation of the user in real-time. In this case, a skeleton can be created and fit to the positions of the marker points. This tends to be slightly less accurate and stable.

A specific technology that is currently quite commonly used in similar situations is the InterSense IS900 (Fig. 7.13). This combines ultrasonic technology with inertial measurements. The IS900 is a common tracker in CAVE™-like installations or other larger spaces, or in situations otherwise not conducive to magnetic tracking.

Other common "cheap" position tracking technologies are marker-based and image-based. The seminal system is the ARToolkit system [14], where the position of a printed marker relative to a camera is derived. On its own, the tracking of a single marker is too noisy to, say, mount a camera on a head to track the head for a HMD display. The marker must be visible for a start. However with sets of markers and calibration of the locations of markers, a versatile position tracker can be built. This type of technology has been commercialized. An example is the InterSense IS1200 system. A novel system using virtual markers in a CAVE™-like environment was proposed by [13]. They proposed placing virtual markers behind the head of the user of the display. These could be used to accurately update the head position, which in turn could be used to move the markers out of the vision of user.

Fig. 7.13 *Left* a typical Intersense IS900 sensor bar installation. *Right* Intersense IS9000 MicroTrax Wand. Courtesy of Intersense, LLC

There are many ways of comparing tracking systems and we refer the interested reader to the previously mentioned reviews [21, 41, 44]. One specific point of comparison that is often overlooked but which is especially important for travel applications is how well the system scales. We can mention at the outset that certain designs do not scale because of the nature of the tracking technology. Magnetic tracking for example does not scale because multiple systems can interfere with each other.

The first wide area tracker of note was the UNC HiBall [42, 43] which is an optical tracker, but which, unlike the optical tracking systems that are now common, is “inside-out,” in that the cameras are mounted on the user pointing towards the ceiling. The ceiling contains strips of LEDs that flash in sequence, and from recognizing points on the ceiling, the position of the camera can be calculated. This technology scales well because there is a single set of cameras, and the LED strips can be mounted over a wide area. Essentially, in the HiBall tracker, the environment is changed to include easily tracked points. In a similar manner, camera-based tracking systems based on markers, such as the InterSense IS1200, can also work over a large area. However the camera arrays are currently bulky, and these inside-out systems can suffer from problems of occlusions of the camera by the body.

Thus, although they are more expensive, large motion capture systems are often used to track large spaces. Motion capture systems such as Vicon are designed to scale up to large numbers of cameras: as the tracked volume increases more cameras are required to retain precision. An impressive system of this type is the VIRTSIM system from Motion Reality Inc. This can track a simple skeleton on 13 users simultaneously over an area of $50'$ by $100'$. It also drives 13 wireless HMDs to give users the impression they are training in a squad with 12 others.

The next class of device is by far the most diverse: hand-held devices. The space of physical and virtual devices that could be used to control travel is vast. It ranges from joysticks through props such as steering wheels, to versatile devices such as smartphones. Many variants have been tried and we will only mention a sampling.

Interested readers are referred to Bowman et al. [5], and the IEEE Symposium on 3D User Interfaces, where new devices are introduced every year.

We start with the simplest device: a button. The very simplest travel interfaces use one or more buttons on a device that the user holds. Indeed, many position-tracking systems include a hand-held device with buttons as standard. A button can be used to indicate that movement should be effected, and the direction of the movement might depend on the orientation of a tracker.

This is all one needs if the display fully surrounds the user (i.e., the field of regard is 360°), but otherwise some control is needed to rotate the viewpoint in the environment. One could achieve this with a button and a gesture to point, but simpler and possibly easier and more intuitive for the user is to provide a joystick (see Fig. 7.12, Right).

There is a broad range of devices that are common with desktop computers and can be used or customized to use in an immersive setup. A wireless handheld gyroscopic mouse is a common choice. As noted previously the gyroscope is only a rotational control, so it is not sufficient to give an accurately registered direction, but a position tracker can be attached.

Assemblies of controls can become quite complicated. Trackers and buttons can be embedded in props such as weapons or sports equipment. While such props have long been common in video games, Hinckley et al. proposed using them for other applications, including embedding a tracker in a doll's head to aid with a neurosurgical visualisation [11]. It may be that the task includes complex controls beyond travel, selection and manipulation of objects. In this case a controller might be custom built or a controller with a wide range of control patterns might be used. For example, the cubic mouse [10] is a tracked cube with three orthogonal rods passing through it.

For more complex user interfaces using multiple dynamic interfaces such as menus, an obvious and common choice is to utilize a mobile device such as a tablet or smartphone as a control device. This can display virtual buttons, virtual sliders, and other controls to facilitate travel. A common use of mobile devices is to act as a secondary data display, by providing maps, e.g., [23], text or numeric information, or alternate views of the virtual world. A metaphor that can be used is the concept of the magic lens or 3D magic lens [39] through which the 3D world can be seen in a different way. This then can act as an interface to tasks in the VE, including travel.

An example of the type of more complex travel interface that can be built with such devices is the World-in-Miniature (WIM) technique, which creates a hand-held miniature map of the world that users can use to move themselves around the VE [26].

Such devices don't need to be physical devices; they can be simulations of physical devices or other representations of virtual controls. For example, the *virtual tricorder* and *pen-and-tablet* approaches embed virtual controls in the 3D environment that are anchored to the tracked position of a hand-held device [4, 35, 46]. These controls can be used for a variety of tasks, including travel.

Our final category of device is hybrid devices. This is a catchall for novel and interesting technologies that can be used to control travel. For travel techniques, an obvious place to put the technology is on the floor. Consumer games utilize devices

Fig. 7.14 Joyman. Courtesy of Julien Pettré, INRIA-Rennes, France

such as dance mats and the Wii Fit Balance Board; these can be made into controllers for VEs very easily. To integrate a larger force sensing surface into a VE requires a larger engineering effort with multiple force sensors. An example of a successful system is the one at McGill University [17]. Because the floor is force sensitive (and in this case includes a haptic display), controls can be integrated into the floor so that the user can step on them. Other novel devices for travel include Joyman (Fig. 7.14) [27], which is based on the concept of a "human-scale joystick." The user stands on a platform that they tilt to control locomotion. A related system is the Virtual Motion Controller (VMC) (Fig. 7.15) [45] which measures the position of a user standing on a plate and then moves the user when they stand on its rim. A variety of foot-based interaction devices exist, such as the Interaction Slippers, which are tracked and sense contact between the feet [16]. Finally, systems such as GAITER (Fig. 7.16) can track foot and leg position and movement but can also support the user's doing actions such as kneeling and going prone [37].

Before moving on to discussing travel techniques themselves, it is worth analysing what these input devices give us in terms of control input. A useful tool is the analysis of Mackinlay et al. [18], which classified interaction devices based on what types of

Fig. 7.15 Virtual Motion Controller. Courtesy of Thomas Furness, Human Interface Technology Laboratory, University of Washington

motion they sense. The type of diagram that they use is shown in Fig. 7.17. It classifies a single device as a union (indicated by connecting dashed lines) of individual sensors (the circles connected with solid lines). We can see that the device might sense linear absolute position (P), linear relative position (dP), linear absolute force (F), linear relative force (dF), rotary absolute position (R), rotary relative position (dR), rotary absolute force (T for tensor), or rotary relative force (dT). We also see how many degrees of freedom each device senses (X, Y, Z for linear degrees of freedom, and rX, rY, rZ for rotary degrees of freedom), and the resolution: from continuous (Inf) through to discrete (1). Thus a common mouse is a combination of a two-dimensional relative position (X, Y) sensor with continuous movement plus typically 2 discrete buttons (Z, but only two position values, indicating 0 or 1) plus a scroll wheel reporting a number of discrete positions of the wheel as relative rotation. The diagram thus depicts the mouse as three separate units (position, buttons, scroll wheel) connected by dotted lines. The two dimensions of the mouse are joined together by a solid line because they are reported by the same sensor in the mouse.

Fig. 7.16 GAITER. Courtesy of Jim Templeman, U.S. Naval Research Laboratory

We have plotted several other common input devices. The Polhemus Fastrak tracker is a widely used magnetic tracker which reports the six-DoF position and orientation of a device in a coordinate system centered on a transmission device. A Wii Remote with MotionPlus is a relatively complex device with four sets of sensors. The first is that the device has 11 buttons on it that (plus an on-off switch that isn't included in the diagram). The second is that the IR camera on the Wii Remote can sense a sensor bar placed on or under the display. This gives effectively four coordinates: three rotations of the Wii Remote relative to the sensor bar, plus a distance of the Wii Remote from the sensor bar. The third sensor is the accelerometer (actually the lowest set of nodes in this connected set), which reports the acceleration of the Wii Remote including gravity. The fourth set is the Wii MotionPlus, which adds gyroscope functionality and thus provides another set of relative rotations in a different coordinate system (the device's coordinate system rather than the screen

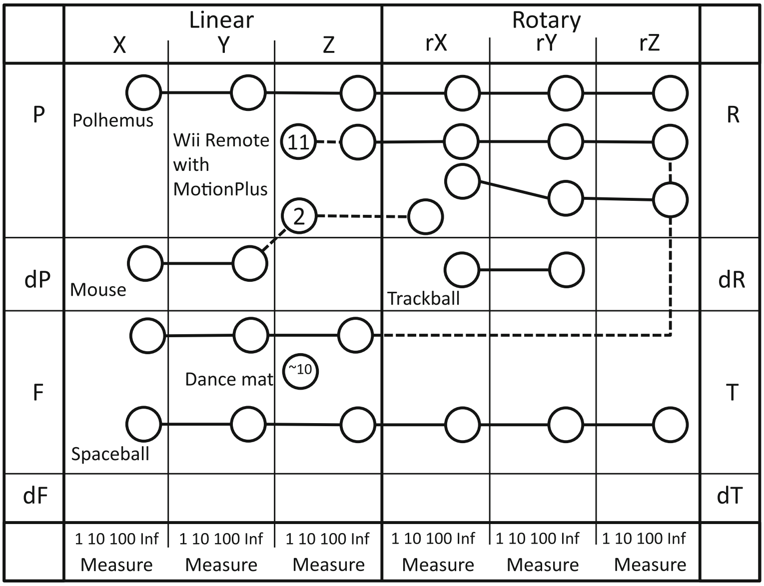

Fig. 7.17 Several input devices plotted using the analysis of Mackinlay et al. [18]

coordinate system). In covering other sections of the diagram we have included some other common devices. A trackball senses two degrees of rotation. A dance mat as might be used for popular dancing games supports approximately 10 buttons that respond to force applied to them, but do not themselves move. Finally the Spaceball is a type of controller that sense forces on it, but does not move. One current example is the HP Spaceball® 5000.

A note: we have not covered treadmills and similar locomotion devices in this chapter. From a control point of view, these provide one or two degrees of input because the user can walk in one or two directions. However the issues in using this for control are subtle. Some treadmill technologies are discussed in Chap. 6 of this volume.

7.4 Travel Techniques

7.4.1 Travel as a Control Task

In the virtual reality model of interaction, to travel the user would walk or use a vehicle, just as he would do in real life. Of course, as already discussed, this ideal is limited by the technology that is available. Most notably, VE systems don't support the user moving themselves over long distances and might even require that the user be stationary. Therefore some sort of virtual travel technique is necessary. Given what we know about sensors and displays, it is first important to identify what we know about the user, and thus what a travel technique actually moves.

We already noted that viewpoint control in the real world is a consequence of movements of various parts of the body. We also note that sometimes the user's head is tracked, and sometimes it is not. If the user stands inside a CAVE™ -like device, then they can physically move a little bit, but may utilize a joystick to move longer distances. Thus in many systems there are effectively two travel techniques: a short-range physical travel technique and a long-range virtual travel technique. This is unlike a standard desktop-style metaphor where the whole of the viewpoint motion control is under a single control metaphor (e.g., a joystick or mouse plus keyboard combination).

In order to understand how these two types of movement work together, we need to describe the various coordinate systems involved, and the relationships among them.

A logical separation which is reified in some implementations is to separate a *display center coordinate* system from a tracker coordinate system. To explain this we need to delve into one potential way of structuring coordinate systems within a display system. An example is shown in Fig. 7.18. This depicts an SSD-type system but the same coordinate systems would usually exist in other display types.

- Tracker Base is the fixed origin of the tracking system coordinate system. Typically this is centered on a physical base unit or some physical component of the tracking system that is static. The tracking unit reports positions relative to this.
- Head Tracker is the relative position of the tracking unit attached to the head. Note that it is not exactly the same as the position of the head, but it is commonly assumed that the tracking unit is in a fixed position relative to the user's head (e.g., on the glasses or on a cap).
- Hand Tracker is the relative position of the tracking unit attached to the hand. Again this is not exactly the same as the hand position as it is assumed there is a known offset between the two.

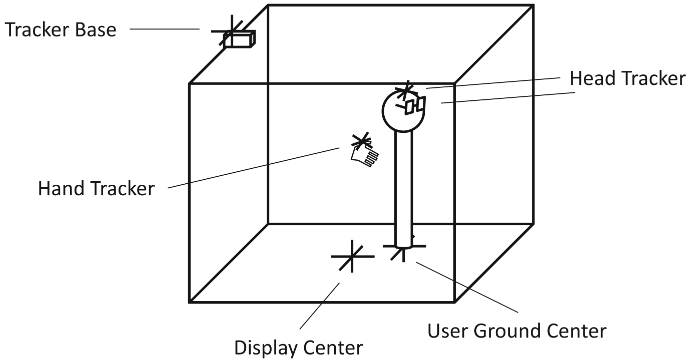

Fig. 7.18 Coordinate systems within a SSD-type display

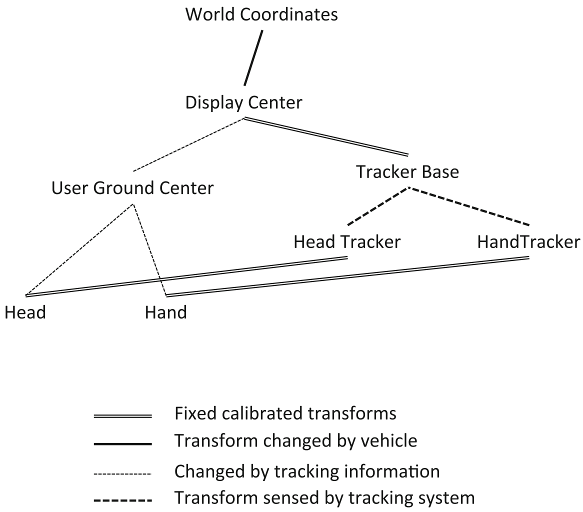

Fig. 7.19 Relationships between coordinate systems in a typical VE system

- Display Center is a fixed position in the display from which measurements of user position are reported. In Fig. 7.18 the display center is the center of the floor.
- User Ground Center is a calculated position relative to the Display Center where the user "is." Typically this is the centroid of his two feet or just the point directly below his head.

The relationships between these coordinate systems are represented in Fig. 7.19, which is an abstract scene graph. The lines represent parent-child relationships between coordinate systems. Each line also represents the transformation from one coordinate system to another. We see that the Display Center is a child of the World Coordinates. The position of the Tracker Base is known relative to the Display Center. This will usually be a one-time calibration at the installation of the tracking system. We can then see that Hand Tracker and Head Tracker coordinates are known relative to the Tracker Base. These are the values reported in real-time by the tracking system. We can thus see that by concatenating the Display Center to Tracker Base and Tracker Base to Hand Tracker transformations, we can find the location of the Hand Tracker relative to the Display Center. The relative position of the Head from the Head Tracker and Hand from the Hand Tracker are also calibrated and fixed. Given that all these transformations are known (the double lines in the figure and the large dashed line), we can then calculate the positions of the Head and Hand relative to the Display Center. However, a common convention, especially when implementing travel techniques is to introduce a User Ground Center coordinate system as defined above. Thus the three light dashed lines must be updated whenever the tracking system reports new positions.

Having explained the roles of coordinate systems, we can see that short-range physical movements are captured in the movements of the Head Tracker and Hand Tracker relative to the Display Center.

For virtual travel over longer distances, the position of the Display Center can be moved relative to the World coordinates. There are many options here. Fundamentally they involves changing the transformation of the Display Center relative to the World origin. This might involve translation, rotation and scale of the Display Center. However, it is very common for certain restrictions to be set. For example, there might be translation only: if the person is surrounded by the display system, they can simply physically turn to see any direction, whereas in a typical SSD with three walls there will need to be a rotation metaphor. If there is translation, then it might be 2D only (i.e., on a plane), or use some form of surface-following algorithm so that the user is always at the same height above the ground. Other questions must be answered as well. If there is translation, in which direction should it be? If there is rotation or scaling, around which point should the scale or rotation occur (e.g., the Display Center, the User Ground Center, or the Head)? And how does the user indicate how to start and stop traveling?

A task decomposition for travel techniques is provided by Bowman et al. [2]. Their decomposition, which is shown in Fig. 7.20, decomposes the travel task into different sub-tasks: start to move, indicate position, indicate orientation, stop moving. To start to move, the user might press and hold a button, and to stop moving she might release

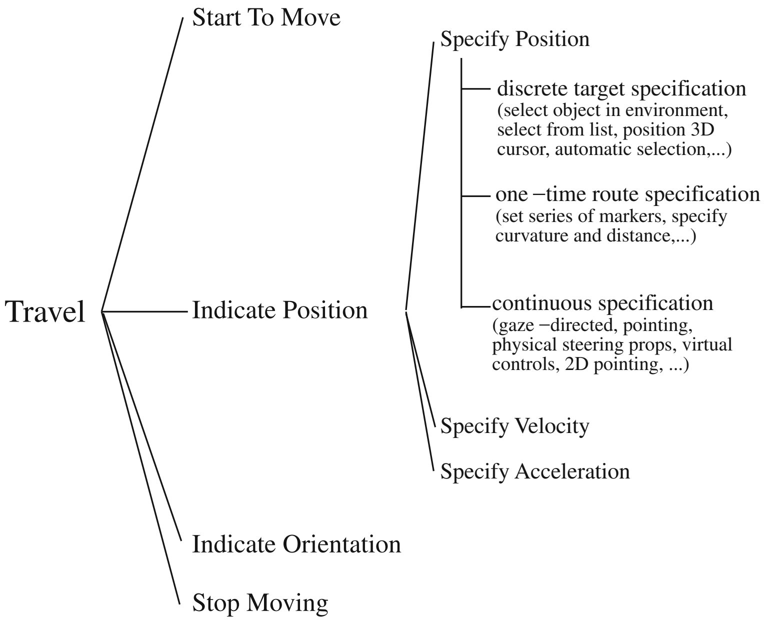

Fig. 7.20 Taxonomy of travel techniques focusing on level of user control [2]

a button. Alternatively, for example, the start and stop might be implicit in the user manipulating a joystick. In our terminology, the indicate position sub-task effects a change in translation of the Display Center, either directly, or by giving the dynamics of change over time (i.e., giving a velocity or acceleration). The task decomposition shows that this can be achieved in three ways: specify position, specify velocity and specify acceleration. Each of these can in turn be achieved in many ways. For example specifying a position can be achieved by selecting a target, giving a route, or by continuous specification. The last of these is the most common in real-time, interactive systems: the user can point or gaze towards the target and effect a control to start and continue travelling. There would be a similar breakdown for orientation: e.g., one might set a target orientation, one might turn using a direct angular control, or one might set a rotation speed.

In looking at the options here we can make a connection to Fig. 7.17, where the input devices had different dimensionality that might map conveniently to the different options here. Obviously a joystick is a common choice for specifying velocity, whereas a mouse might be better deployed for relative rotation or clicking to select targets in the world.

We examine some common control configurations in the following sections. The reader is also referred to Bowman et al. [5].

7.4.2 *Direct Self Motion Control Techniques*

In this section we cover direct control, where the user has continuous control over the direction of travel from a first-person view at every update cycle of the simulation.

The most obvious technique for travel is gaze-directed steering [22]. This is the default travel technique in many immersive systems, and it is also found in many 3D games. When the user makes the relevant control input (e.g., presses a button or moves the joystick forwards), the Display Center coordinate system moves forward along the direction of gaze. In a desktop virtual reality system, this would be through the center of the screen. In an immersive VE system, this is typically in the direction of a line in the center of the two eye lines. There are many variants depending on the control input as suggested previously. The velocity of travel might be constant, a joystick deflection might control the velocity of travel, or it might set an acceleration. The movement might ease in or ease out when the control is changed. If a joystick is used, typically it is forward/backward that controls travel along the direction of travel, but the other axis might map to strafing (sideways movement) or turning (rotation), or it might be ignored. Finally, the actual direction of travel might be clamped to be in only two dimensions or so that the user is at a set height above the ground.

Gaze-directed steering has the advantage that it is simple to explain to users, but it has the major disadvantage that the user must look in the direction they intend to travel. There are obvious variants: one could use any other coordinate system or relation between tracked points to set the travel direction. The most obvious ones are flying in the direction of pointing with a hand tracker [22], direction of gaze recorded by an eye-tracker, direction of the torso, the relative direction of hand from eye or the relative position of the two hands. All of these have been implemented in VE systems. Pointing with a hand tracker is a very commonly implemented technique because it decouples head orientation from travel so that users can look around while moving [2].

7.4.3 Indirect Self Motion Control Techniques

Indirect control techniques set the target of travel in an indirect or asynchronous manner. The most common technique is to indicate a target in the environment, and then to enable travel to that position over a period of time or instantly (when it is known as teleporting). Indicating the target might involve simply targeting an object by pointing directly (e.g., "go over there"), or on a map or miniature world. An example of the former is the ZoomBack technique [47] that uses a typical ray-casting metaphor to select an object in the environment, and then moves the user to a position directly in front of this object. Ray-casting has been used in other 3D interfaces for target-based travel as well, e.g., [3]. An example of the latter is the previously mentioned WIM technique [26], in which a small human figure represents the user's position and orientation in the miniature world. The user places the user representation in the WIM, and then a path is calculated that moves the camera to this location, taking into account any rotation that is necessary to reach the target orientation. In this technique the transition could be achieved by zooming in to the WIM itself, or it could be planned as a motion through the VE at the original scale.

7.4.4 Scene Motion Techniques

A common alternative to self-motion travel is to manipulate the scene. This is very common on desktop interfaces where the metaphor is that the camera is static and the object on the screen is moving. Less obvious is that this can be turned in to a travel technique, in that the object is considered stationary and the camera is moved around the object. The rotations of the camera for an immersive system would be violent, so this is not very commonly done, but the translation equivalent (pull and push objects) has been demonstrated. Most notable is the "grab the air" technique [19, 40]. In this technique the user can grab anywhere in the environment and when they move her hand back and forth, they move themselves through the environment. This motion can be scaled by using two hands. Similar techniques include Mine's Scaled-World Grab technique [24] and the LaserGrab technique of Zeleznik et al. [47]. A related set of techniques involve manipulation of objects using image plane techniques [28].

7.4.5 Other Control Inputs

The interaction techniques for travel that we have discussed in the section above cover a broad range of those that are used in practice. However, the area has seen a number of innovative techniques. Another way of controlling speed of rotation or velocity of travel involves measuring the distance between points on the body and using that as the rate. For example, the distance between the head and a measured or nominal foot position gives an estimate of lean, and this can be used to control velocity [7, 16, 32]. Alternatively the distance between hand and head can be used to control velocity in a point to fly technique [22].

An alternative to using a device to effect travel is to track a user movement that is similar to walking. Such techniques are called "walking in place" metaphors, where users move their feet to simulate walking without actually translating their bodies [31]. In the case of Slater et al., the user had to mimic walking, and a gesture recognition system detected that the user was performing this mime by monitoring his head movement. Walking in place metaphors have attracted a lot of interest because users have to physically exert themselves. A novel platform that allowed walking in place with extended leg movement was presented by Swapp et al. [34]. Walking in place techniques are covered in Chaps. 10 & 11.

Finally, we note that especially with four-walled CAVE™-like systems and HMDs with restricted tracker spaces, there has been a lot of interest in techniques that bias rotation to achieve the effect that the user doesn't look away from the main walls, or walks in the correct direction. These are covered elsewhere in the book (see Chap. 14), but we note the work by Razzaque et al. on redirected walking [29] which provides imperceptible rotation distortion, and explicit amplification of rotation [16].

7.5 Conclusion

We hope that in this short introduction to interaction devices and displays for virtual walking, we have conveyed some of the challenges of the field and the constant innovation that there has been over the past couple of decades. Travel is a very hard problem for virtual reality systems: it is inherently a "two-task" system because the user can move physically for short maneuvering tasks, but the user also needs a virtual travel technique to move over long distances. Other chapters in this volume indicate some of the work that is being done to alleviate the need for a virtual travel technique in some situations, but virtual travel techniques will likely be with us for some time yet. In describing the range of different display types, interaction devices and control methods, it should also be obvious that there are no one-size-fits-all solutions and that techniques need customization to fit application needs and system capabilities. While there are many options, there are also established best practices that can be uncovered by studying the literature. There will no doubt be further innovation, especially in gesture-based control that is enabled by recent advances in camera technology. We look forward to testing new techniques as they emerge.

References

1. Agrawala M, Beers AC, McDowall I, Fröhlich B, Bolas M, Hanrahan R (1997) The two-user Responsive Workbench: support for collaboration through individual views of a shared space. In: Proceedings of the 24th annual conference on computer graphics and interactive techniques (SIGGRAPH '97). ACM Press/Addison-Wesley, New York, pp 327–332
2. Bowman D, Davis E, Badre A, Hodges L (1999) Maintaining spatial orientation during travel in an immersive virtual environment. Presence: Teleoper Virt Environ 8(6):618–631
3. Bowman D, Johnson D, Hodges L (2001) Testbed evaluation of VE interaction techniques. Presence: Teleoper Virt Environ 10(1):75–95
4. Bowman D, Wineman J, Hodges L, Allison D (1998) Designing animal habitats within an immersive VE. IEEE Comp Graph Appl 18(5):9–13
5. Bowman D, Kruijff E, LaViola J, Poupyrev I (2004) 3D user interfaces: theory and practice. Addison Wesley Longman, Redwood City
6. Cruz-Neira C, Sandin DJ, DeFanti TA (1993) Surround-screen projection-based virtual reality: the design and implementation of the CAVE. In: Proceedings of the 20th annual conference on computer graphics and interactive techniques (SIGGRAPH '93). ACM, New York, pp 135–142
7. Fairchild K, Hai L, Loo J, Hern N, Serra L (1993) The heaven and earth virtual reality: designing applications for novice users. In: Proceedings of IEEE virtual reality annual international symposium, 1993. IEEE, pp 47–53
8. Feiner S, MacIntyre B, Höllerer T, Webster T (1997) A touring machine: prototyping 3D mobile augmented reality systems for exploring the urban environment. In: Proceedings of first international symposium on wearable computers. IEEE, pp 74–81
9. Fernandes KJ, Raja V, Eyre J (2003) Cybersphere: the fully immersive spherical projection system. Commun. ACM 46(9):141–146
10. Fröhlich B, Plate J (2000) The cubic mouse: a new device for three-dimensional input. In: Proceedings of the SIGCHI conference on human factors in computing systems (CHI '00). ACM, New York, pp 526–531
11. Hinckley K, Pausch R, Goble JC, Kassell NF (1994) Passive real-world interface props for neurosurgical visualization. In: Proceedings of ACM CHI'94 conference on human factors in computing systems. ACM New York, pp 452–458
12. Höllerer T, Kuchera-Morin J, Amatriain X, (2007) The allosphere: a large-scale immersive surround-view instrument. In: Proceedings of the 2007 workshop on emerging displays technologies: images and beyond: the future of displays and interacton (EDT '07). ACM, New York, Article 3
13. Hutson M, Reiners D (2011) JanusVF: accurate navigation using SCAAT and virtual fiducials. IEEE Trans Visual Comput Graph 17(1):3–13
14. Kato H, Billinghurst M (1999) Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In: Proceedings of the 2nd IEEE and ACM international workshop on augmented reality (IWAR 99), pp 85–94
15. Kulik A, Kunert A, Beck S, Reichel R, Blach R, Zink A, Fröhlich B (2011) C1x6: a stereoscopic six-user display for co-located collaboration in shared virtual environments. In: Proceedings of the 2011 SIGGRAPH Asia conference (SA '11). ACM, New York, Article 188, 12 pp
16. LaViola J, Acevedo D, Keefe D, Zeleznik R (2001) Hands-free multi-scale navigation in virtual environments. In: Proceedings of the 2001 symposium on interactive 3D graphics (I3D '01). ACM, New York, pp 9–15
17. Law AW, Peck B, Visell Y, Kry P, Cooperstock J (2008) A multi-modal floor-space for experiencing material deformation underfoot in virtual reality. In: Proceedings of IEEE international workshop on haptic audio visual environments and games. IEEE, pp 126–131
18. Mackinlay J, Card SK, Robertson GG (1990) A semantic analysis of the design space of input devices. Human-Comput Inter 5(2):145–190
19. Mapes D, Moshell J (1995) A two-handed interface for object manipulation in virtual environments. Presence: Teleoper Virt Environ 4(4):403–416

20. McMahan R, Bowman D, Zielinski D, Brady R (2012) Evaluating display fidelity and interaction fidelity in a virtual reality game. IEEE Trans Visual Comput Graph (Proc IEEE Virt Real) 18(4):626–633
21. Meyer K, Applewhite H (1992) A survey of position trackers. Presence: Teleoper Virt Environ 1(2):173–200
22. Mine M (1995) Virtual environment interaction techniques. UNC Chapel Hill CS Dept
23. Miguel MM, Ogawa T, Kiyokawa K, Takemura H (2007) A PDA-based see-through interface within an immersive environment. In: Proceedings of The 17th international conference on artificial reality and telexistence (ICAT 2007), pp 113–118
24. Mine M, Brooks F, Sequin C (1997) Moving objects in space: exploiting proprioception in virtual environment interaction. In: Proceedings of the 24th annual conference on computer graphics and interactive techniques (SIGGRAPH '97). ACM Press/Addison-Wesley, New York, pp 19–26
25. Office of Naval Research (2012) Infantry immersion trainer. http://www.onr.navy.mil/Media-Center/Fact-Sheets/infantry-immersion-trainer.aspx, Retrieved July 2012
26. Pausch, R., Burnette, T., Brockway, D., & Weiblen, M. (1995). Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniature. In Proceedings of the 22nd annual conference on Computer graphics and interactive techniques (SIGGRAPH '95). ACM, New York, pp 399–400
27. Pettré J, Siret O, Marchal M, de la Rivire J-B, Lécuyer A (2011) Joyman: an immersive and entertaining interface for virtual locomotion. In SIGGRAPH Asia 2011 emerging technologies (SA '11). ACM, New York, Article 22
28. Pierce J, Forsberg A, Conway M, Hong S, Zeleznik R, Mine M (1997) Image plane interaction techniques in 3D immersive environments. In: Proceedings of the 1997 symposium on interactive 3D graphics (I3D '97). ACM, New York, 39-ff
29. Razzaque S, Swapp D, Slater M, Whitton M, Steed A (2002) Redirected walking in place. In: Proceedings of the workshop on virtual environments 2002 (EGVE '02), Eurographics Association, Aire-la-Ville, Switzerland, pp 123–130
30. Slater M, Steed A, Chrysanthou Y (2001) Computer graphics and virtual environments: from realism to real-time. Addison-Wesley, New York
31. Slater M, Usoh M, Steed A (1995) Taking steps: the influence of a walking technique on presence in virtual reality. ACM Trans Comput-Human Inter 2(3):201–219
32. Song D, Norman M (1993) Nonlinear interactive motion control techniques for virtual space navigation. In: Proceedings of IEEE virtual reality annual international symposium. IEEE, pp 111–117
33. Sutherland IE (1965) The ultimate display. Proc IFIP 65(2):506–508
34. Swapp D, Williams J, Steed A (2010) The implementation of a novel walking interface within an immersive display. In: Proceedings 2010. IEEE symposium on 3D user interfaces (3DUI). IEEE, pp 71–74
35. Szalavári Z, Gervautz M (1997) The personal interaction panel—a two-handed interface for augmented reality. Comput Graph Forum (Proceedings of EUROGRAPHICS'97. Budapest, Hungary) 16(3):335–346
36. Tamura H, Yamamoto H, Katayama A (2001) Mixed reality: future dreams seen at the border between real and virtual worlds. IEEE Comput Graph Appl 21(6):64–70
37. Templeman JN, Sibert LE, Page RC, Denbrook PS (2009) Designing user interfaces for training dismounted infantry. In: Nicholson D, Schmorrow D, Cohn J (eds) Handbook of virtual environments for training and education, vol 2. Greenwood Publishing Group, Westport, pp 219–239
38. Usoh M, Arthur K, Whitton M, Bastos R, Steed A, Slater M, Brooks F (1999) Walking > Walking-in-place > flying. In: Virtual environments. Proceedings of the 26th annual conference on computer graphics and interactive techniques (SIGGRAPH '99). ACM Press/Addison-Wesley, New York, pp 359–364
39. Viega J, Conway MJ, Williams G, Pausch R (1996) 3D magic lenses. In: Proceedings of the 9th annual ACM symposium on user interface software and technology (UIST '96). ACM, New York, pp 51–58

40. Ware C, Osborne S (1990) Exploration and virtual camera control in virtual three dimensional environments. In: Proceedings of the 1990 symposium on interactive 3D graphics (I3D '90). ACM, New York, pp 175–183
41. Welch G (2002) Tracking bibliography. http://www.cs.unc.edu/~tracker/ref/biblio/index.html. Retrieved July 2012
42. Welch G, Bishop G, Vicci L, Brumback S, Keller K, Colucci D (1999) The HiBall tracker: high-performance wide-area tracking for virtual and augmented environments. In: Proceedings of the ACM symposium on virtual reality software and technology (VRST '99). ACM, New York, pp 1–10
43. Welch G, Bishop G, Vicci L, Brumback S, Keller K, Colucci D (2001) High-performance wide-area optical tracking. Presence: Teleoper Virt Environ 10(1):1–21
44. Welch G, Foxlin E (2002) Motion tracking: no silver bullet, but a respectable arsenal. IEEE Comput Graph Appl 22(6):24–38
45. Wells M, Peterson B, Aten J (1996) The virtual motion controller: a sufficient-motion walking simulator. In: Proceedings of IEEE virtual reality annual international symposium. IEEE, pp 1–8
46. Wloka M, Greenfield E (1995) The virtual tricorder: a uniform interface for virtual reality. In: Proceedings of the 8th annual ACM symposium on User interface and software technology (UIST '95). ACM, New York, pp 39–40
47. Zeleznik R, LaViola J, Feliz D, Keefe D (2002) Pop through button devices for VE navigation and interaction. In: Proceedings of IEEE virtual reality 2002. IEEE, pp 127–134