

Roy A. Ruddle

royr@comp.leeds.ac.uk

The Informatics Research Institute

School of Computing

University of Leeds

Leeds LS2 9JT UK

Dylan M. Jones

School of Psychology

Cardiff University

Cardiff CF10 3YG UK

Movement in Cluttered Virtual Environments

Abstract

Imagine walking around a cluttered room but then having little idea of where you have traveled. This frequently happens when people move around small virtual environments (VEs), searching for targets. In three experiments, participants searched small-scale VEs using different movement interfaces, collision response algorithms, and fields of view. Participants' searches were most efficient in terms of distance traveled, time taken, and path followed when the simplest form of movement (view direction) was used in conjunction with a response algorithm that guided ("slipped") them around obstacles when collisions occurred. Unexpectedly, and in both immersive and desktop VEs, participants often had great difficulty finding the targets, despite the fact that participants could see the whole VE if they stood in one place and turned around. Thus, the trivial real-world task used in the present study highlights a basic problem with current VE systems.

I Introduction

An important class of virtual environment (VE) application is one that uses a small but cluttered environmental setting. Good examples are applications that are used to assess ergonomic aspects of work places, or the maintainability of large industrial equipment. A primary goal of these applications is to allow users to move around and interact with objects in the VE in the same way as they would in the real world. If this can be achieved, then detailed and realistic human-in-the-loop design assessments can be performed using 3-D CAD/CAM (virtual) prototypes long before products are physically built, simultaneously reducing costs, reducing time to market and improving product quality.

This article is concerned with one aspect of user interaction in cluttered VEs: the problem of how users move themselves around. Many different interface devices have been developed (Hand, 1997; Templeman, Denbrook, & Sibert, 1999), and a variety of these have been used to allow people to move around large-scale VEs (Weatherford, 1985) in studies of navigation (Chance, Gaunet, Beall, & Loomis, 1998; Darken & Sibert, 1996; Ruddle, Payne, & Jones, 1999b; Witmer, Bailey, Knerr, & Parsons, 1996). However, a characteristic of these VEs was that the task of movement was straightforward, typically one of traveling down corridors in a virtual building or across a virtual sea. There is a dearth of behavioral research into the use of any interface for movement within a cluttered VE, in which the task of movement is complicated by the need to avoid obstacles.

Three important characteristics of interfaces for cluttered VEs are the type of

movement that is allowed, the device that is used to achieve that movement, and what happens when the user collides with an object. These are discussed in the following sections, together with their likely effect on the task used in the experiment.

1.1 Movement and Devices

Types of movement are best compared by considering movement in the real world. When a person travels in the real world, three types of directional movement in the horizontal plane are important. First, there is the orientation (heading) of the person's body (H_b). Then there are the person's direction of view (H_v —they can look around while moving) and travel (H_t —they can move in any direction, such as forwards or sideways). In addition, the person can vary his or her velocity in both a positive and negative direction.

Most VE interfaces use view-direction travel (travel where you look; $H_t = H_b = H_v$) (Darken & Sibert, 1996; Witmer et al., 1996) or body-direction travel (sometimes called hand-direction travel), in which the view and travel directions are decoupled ($H_t = H_b \neq H_v$; the \neq sign is used to indicate that two directions can be varied independently) (Bowman, Koller, & Hodges, 1997; Ruddle et al., 1999b). With both of these, users have to turn before they can travel in a new direction, and an additional restriction is that most implementations have only allowed forward (positive) movement. Occasionally, independent movement has been implemented, allowing users to move directly in any direction relative to the orientation of their bodies ($H_t \neq H_b \neq H_v$) (Chance et al., 1998).

All three types of movement can be implemented in both immersive and desktop VEs using a variety of different devices, and, clearly, the device that is chosen can have a substantial effect on user performance. Interface devices can be characterized by factors such as the number of degrees of freedom (DOFs) that can be simultaneously varied, the order of control that is used (zero-order (position), first-order (velocity) or second-order (acceleration)), and the range of values that are mea-

sured. (For a review, see Card, Mackinlay, and Robertson (1991).)

In immersive VEs, sensors are used to track the movement of parts of a user's body and the user typically controls his or her speed by pressing or holding down buttons. The use of two sensors (such as head and waist) allows body-direction movement, whereas the use of just one sensor generally restricts movement to taking place in the user's direction of view, although an exception was the interface implemented by Ruddle et al. (1999b). The devices used to allow independent movement range from those in which a user is physically positioned in a large, empty room and actually walks round a VE, walks on a 2-D treadmill, or walks in place (Templeman et al., 1999) to devices that provide the same DOFs of movement but which the person uses with a different set of musculature (for example, the use of a joystick).

In desktop VEs, a mouse and keyboard are the most common interface devices, and these can be used to implement either view- or body-direction movement. Examples of the latter that allow movement along straight and curved paths can be found in Ruddle, Payne, and Jones (1997, 1999a), respectively. A joystick or cursor keys are suitable for implementing independent movement on the desktop.

1.2 Collision Detection and Response

In the real world, people (usually) walk around obstacles with little conscious effort. In VEs, users often cannot see the obstacles they are about to bump into because the field of view (FOV) is limited, and the lack of fine movement control means that obstacles are difficult to avoid at the last moment. However, in VEs used for applications such as ergonomic design, collision detection is important so that realism is maintained and obstacles are impediments to movement, rather than just being a source of visual clutter. Collision response (what happens after a collision has occurred) can be performed in a variety of ways. The simplest of these is to prevent the user from moving to the colliding position (a stop algorithm), and this, of course, is what happens in the real world. Other responses guide the user

around obstacles, with common versions being slip (Jacobson & Lewis, 1997) and force-field (Xiao & Hubbold, 1998) algorithms.

With guidance algorithms, the VE software performs the fine detail of navigational movements, and users only have to indicate the general direction in which they would like to travel, meaning that the process of movement becomes somewhat different to the real world. Therefore, when considering the design of VE applications that will be used to train for or analyze real-world situations, it is important to make a distinction between collisions that occur because of the amount of clutter that is present (that is, collisions that are likely to occur in the real world) and collisions that occur primarily because of limitations of the interface. For the former, a stop response algorithm should be implemented, but a guidance algorithm is more appropriate for the latter.

1.3 Experimental Task

To investigate movement in cluttered VEs, we designed a task that was equivalent to a person walking around a cluttered room looking for target objects in known, possible locations. The participants' goal was to find the targets, so the interface was simply something that the participants used to achieve their goal, and this reflects the role of interfaces in VE applications.

To perform the experimental task in an efficient manner, participants had to travel through the VE while remembering where they had already searched. This involved the participants in establishing an effective frame of reference and keeping track of changes in their position and orientation. The simpler the interface was to use, the greater the amount of cognitive resources that the participants could devote to updating and maintaining their spatial memory, but increases in the amount that the participants had to turn around produced corresponding increases in the amount of spatial information that they needed to process.

The simplicity of an interface is affected by the mapping between the physical movements that the participant performs and the movement they make in a VE and the number of DOFs that are being controlled. When view- or body-direction movement is performed

in an immersive VE, sensors are used to measure participants' physical changes of orientation, and corresponding changes are made to their orientation in the VE (zero-order control). However, participants change their position by controlling their speed (first- or second-order control) rather than by direct physical movement. When participants walk to achieve independent movement, they use zero-order control to change their position and orientation. However, with a device such as a joystick, orientation changes are controlled as for view- and body-direction movement (that is, zero-order), but position changes take place using a speed of movement that is related to the changes in the joystick's position (that is, first- or second-order). Zero-order control allows more precise movement than the other orders of control, and this means that devices that provide zero-order control are likely to require fewer cognitive resources to use.

All of the devices used in the present study provided zero- and first-order control for changes in orientation and position, respectively. In general, interfaces that have fewer DOFs require fewer cognitive resources to control, and this means that, in the present study, view-direction movement required fewer resources than did body-direction or independent movement.

To provide a frame of reference, each of the four walls of the VEs was a different color. With body-direction and independent movement, participants could look around while traveling and, in the experiments, select targets without changing their direction of travel, thereby reducing the amount of work the participants had to perform during spatial updating. When view- or body-direction movement was being used with the stop response algorithm, the participants had to turn around to move away from an obstacle each time they collided (they could move only forwards, not backwards) and turn again to continue traveling in the same direction as before. However, when independent movement was being used, the participant could simply side-step, again reducing the amount of work that the participant had to perform during spatial updating.

Thus, on one hand, simple interfaces such as view-direction movement reduce the effort needed for control, but, on the other hand, they increase the work that

Table 1. *The Display Used and Factors Investigated in Each Experiment*

Experiment	Display	Factors investigated	Factors inherited
1	HMD	View- vs. body-direction movement Narrow vs. wide gaps between obstacles	—
2	HMD	View-direction vs. independent movement Slip vs. stop collision response algorithm	Narrow VE
3	Desktop	Normal vs. wide FOV	Narrow VE, View-direction movement, Slip collision response algorithm

NOTE. Factors marked as inherited are those that were used in all conditions of an experiment, as a result of earlier findings.

a participant has to perform during spatial updating. With training, even moderately complex interfaces become largely automatic to use, meaning that the issue of spatial updating is likely to be of primary importance.

This interaction between interface simplicity and spatial updating means that it is not possible to predict which type of movement will be most effective for cluttered VEs; indeed, different types of movement may be differentially affected by different amounts of clutter. The remainder of this article describes three experiments that investigated different types of movement and collision response algorithms in cluttered VEs. The factors investigated in each experiment are shown in Table 1.

2 Experiment I

Participants wore a head-mounted display (HMD) and used view- and body-direction interfaces to search two environments that had either narrow or wide gaps between obstacles (a large or small amount of clutter, respectively). A repeated-measures design was used, with every participant performing searches in each of the four conditions. The dependent variables were distance traveled, time taken, number of collisions, and path efficiency.

2.1 Method

2.1.1 Participants. Twelve participants (eight men and four women) took part in the experiment, and their ages ranged from 19 to 32 years. All the participants were either undergraduates or graduates who volunteered for the experiment and were paid an honorarium for their participation. One participant withdrew due to nausea and was replaced in the experiment.

2.1.2 Materials. The VE software was a C++ Performer application that was designed and programmed by the authors and ran on a Silicon Graphics Maximum IMPACT workstation. The HMD was a Virtual Research VR4 (247×230 pixel resolution, 48×36 deg. FOV), and head-tracking was performed using a Polhemus FASTRAK sensor and the MR Toolkit (Green, 1995). Images were displayed in stereo in the HMD, the interpupillary distance was adjusted for each participant, and the application update rate was 12.5 Hz.

The environments (see figures 1 and 2) were walled enclosures that contained 33 cylinders. The wall was 1.5 m high, and its sides were different colors. The cylinders were all 0.5 m in diameter. Their layout comprised eight identical groups of four, with the 33rd cylinder positioned in the center. In the narrow VE, the distance between the wall and the outsides of the near-




Figure 1. Plan view of the narrow (left) and wide (right) environment.

est cylinders, and the smallest distance between the outsides of any two cylinders, was 0.75 m. The minimum distance between the outsides of any two cylinders in each group was 1.268 m. In the wide VE, these two distances were 1.750 m and 2.682 m, respectively. Although the minimum gap between obstacles in the two environments differed by only 1 m, the clearance between the obstacles and the participant's "body" (a cylinder that was 0.5 m in diameter and 0.2 m high) when they traveled through that gap varied by a factor of five (0.25 m versus 1.25 m). For each trial, the central cylinder and two of the cylinders in each group were 1.35 m high and colored green. The other two cylinders in each group were 1.5 m high and two-colored. (The bottom 1.35 m was green, and the remainder was cyan.) Inside one of these was a target (a 0.1 m by 0.1 m white square), but the other was empty (a decoy). Within each group, the cylinders that were two-colored were chosen randomly for each trial.

Participants moved around the VEs while physically standing up. When using the view-direction interface, participants always traveled in the direction in which they were looking, as measured by the FASTRAK sensor that was attached to the HMD. When using the body-direction interface, participants wore a second FASTRAK sensor on a belt around their waist. Their direction of view was decoupled from their direction of movement, with the latter measured by the waist sensor.




Figure 2. Interior view of the narrow VE.

With both interfaces, only the heading and pitch data from the HMD sensor was used. (Roll and translational movements were ignored.) Participants held a pistol grip in one hand and used one button on the grip to control their speed and another button to select the targets. If the participant held down the speed button, they accelerated at 0.75 m/s^2 to a maximum speed of 1.5 m/s (walking pace) but stopped as soon as they released the button. (It acted as a "dead man's handle".) All participants were given the same virtual eye height (1.65 m) to ensure that the targets were always visible from the same distance, irrespective of a participant's actual height. With this virtual eye height, the center of each target was visible from a distance of 0.747 m. Collision detection, implemented using the RAPID software library (Gottschalk, Lin, & Manocha, 1996), meant that participants' translational movements were stopped if their body cylinder collided with the wall or any of the one- or two-colored cylinders. Participants could move again only if they traveled in a noncolliding direction.

2.1.3 Procedure. Participants were run individually and took approximately 2.5 hrs. to complete the experiment. A participant first practiced moving around the narrow VE with one of the interfaces

(such as view-direction) for a total of 30 min., structured as six sessions lasting 5 min. each. During this time, the participant was also familiarized with the searching task. Then the participant underwent four trials in one of the environments (for example, wide) and four trials in the other environment (for example, narrow). Following that, the participant practiced the other interface (three sessions of 5 min. each; less practice was required because the participant was already familiar with the task) and did four trials in each environment. The order in which the interfaces and environments were used was balanced using a Latin square design. For each trial, participants started in the wall's recess and searched the environment until they had found and selected all eight targets. Participants were informed that the targets were always in the two-colored cylinders, but that the position of those cylinders changed between trials. As a precautionary measure, symptoms of VE sickness were monitored for one hour at the end of the experiment, using the Short Symptom Checklist (SSC) developed by Cobb, Nichols, Ramsey, and Wilson (1999), but these data are not reported here.

2.2 Results

Participants' performance in each trial was measured using four primary types of data: distance traveled, time taken, number of collisions, and the path efficiency. The distributions of all three types were normalized using a logarithmic transformation. Initial analyses showed no effect of trial number, so the data reported here were analyzed using two-factor (interface \times width of the gap between obstacles), repeated-measures analyses of variance (ANOVAs).

As expected, participants traveled approximately half as far in the narrow environment ($M = 47.8$ m, $SD = 8.8$) as in the wide environment ($M = 93.2$ m, $SD = 26.3$) ($F(1, 11) = 742.19$, $p < 0.01$). However, there was little difference between the distances participants traveled when using the view-direction ($M = 63.9$ m, $SD = 32.9$) and body-direction interfaces ($M = 62.5$ m, $SD = 42.5$) ($F(1, 11) = 0.38$, $p > 0.05$).

Three different types of analysis were performed using




Figure 3. Mean time spent moving and stationary (excluding time spent selecting the targets) for each interface and VE in experiment 1. Error bars indicate standard error of the mean.

the time data: total trial time, time spent stationary (not translating), and time spent moving. The time that participants spent stationary while selecting each target was excluded. The total trial time was the sum of the stationary and moving times, and all three are shown in figure 3. For the total trial time, participants performed the trials significantly faster with the view-direction interface than with the body-direction interface ($F(1, 11) = 8.18$, $p < 0.05$), but there was no effect of gap width ($F(1, 11) = 4.40$, $p > 0.05$). Similarly, participants spent significantly less time stationary with the view-direction interface than with the body-direction interface ($F(1, 11) = 16.25$, $p < 0.01$), but there was no effect of gap ($F(1, 11) = 3.92$, $p > 0.05$). However, for the time spent moving, there was no effect of interface ($F(1, 11) = 1.59$, $p > 0.05$), but participants spent significantly less time moving in the narrow VE than in the wide VE ($F(1, 11) = 54.52$, $p < 0.01$).

For the purposes of analysis, the number of distinct collisions in each trial was calculated. To avoid multiple counts when a participant collided with the same obstacle several times in quick succession, the number of collisions was incremented only if the participant had moved at least 0.5 m (the diameter of their body cylinder) since the last collision. Collisions that occurred while selecting or leaving a target were excluded. Partic-




Figure 4. Mean number of collisions for each interface and VE in experiment 1. Error bars indicate standard error of the mean.

Participants collided more often in the narrow VE than in the wide VE ($F(1, 11) = 28.23, p < 0.01$), but there was no effect of interface ($F(1, 11) = 0.64, p > 0.05$). See figure 4.

In each trial, path efficiency was measured as the extent to which participants retraced their steps. This was determined by calculating the number of two-colored cylinders (the possible target locations) that participants visited more than once. For a visit to count, the center of a cylinder (the location of the target square, if there was one) had to lie within a participant's FOV. A subsequent visit wasn't counted until the participant moved closer to another two-colored cylinder and further away from the original cylinder than the minimum target detection distance. (Either criterion could be met in the narrow VE without the other one being satisfied.)

The percentage of two-colored cylinders that were revisited did not differ significantly between the two interfaces ($F(1, 11) = 0, p > 0.05$), or the two gaps ($F(1, 11) = 2.76, p > 0.05$). Means (standard deviations) for the four conditions were view-narrow (21.4% (12.4)), view-wide (12.9% (17.5)), body-narrow (17.4% (13.6)), and body-wide (16.1% (16.3)). This, however, does not tell the whole story. In 22% of the trials, participants revisited half (or more) of these cylinders, and each participant did this at least once. Sometimes a contributory cause was the fact that a participant had al-




Figure 5. Least efficient path (largest number of two-colored cylinders revisited) for each condition in experiment 1. Clockwise, from top left, the conditions are view-narrow, body-narrow, body-wide, and view-wide path. The one- and two-colored cylinders are shown as wireline and solid circles, respectively. Each path was taken by a different participant.

ready traveled past a target without turning to look at it, but on other occasions they repeatedly traveled through some parts of the VE while completely neglecting others. (See figure 5.)

2.3 Discussion

Participants' searches were expected to be quicker with the body-direction interface because they could look around while traveling, but the opposite effect occurred and they performed the trials quickest with the view-direction interface. In each condition, participants spent a large proportion of the time stationary, either looking around to decide in which direction to travel or recovering from a collision, and it was this time that provided most of the difference between the two inter-

faces. With both interfaces, participants collided with obstacles more times and more often in the narrow VE than in the wide VE.

The most unexpected finding of the experiment was the frequency with and extent to which participants retraced their steps. With both interfaces and both environments, participants became easily disoriented and often seemed to have little idea of where they had traveled. The difficulties they experienced were particularly unexpected for four reasons.

First, the environments were small-scale spaces (Weatherford, 1985) that participants could see the whole of by standing in one place and turning around. Second, participants could perform the task efficiently if they traveled systematically through the VEs and looked on top of each two-colored cylinder. Third, the four walls of the VEs were each a different color, providing a visual frame of reference. Fourth, participants' changes of direction in the VEs corresponded with their physical changes of direction, which is something that has been shown to significantly help people maintain their orientation in the real world (Presson & Montello, 1994; Rieser, 1989).

One important factor that contributed to the inefficiency of participants' searches was the frequency with which they collided with the wall and cylinders, and this occurred more frequently as the amount of clutter increased. Each time participants collided, the VE software stopped them from moving; to start again, they had to turn around and travel away from the colliding object. As well as slowing participants down, this probably had a disorienting effect because it forced them to change direction much more than they would in the real world, where they could simply sidestep (independent movement).

Two methods can be proposed for improving the ease and efficiency with which people can move around cluttered VEs. The first method is to implement independent movement, and the second method is to implement a collision response algorithm that guides participants around obstacles so that changes in their view direction are reduced. Experiment 2 investigated both of these methods.

3 Experiment 2

In experiment 2, participants searched the narrow VE using both view-direction and independent movement interfaces, and stop and slip collision response algorithms. Independent movement was implemented using a virtual joystick. Two predictions were made. First, the slip algorithm was expected to reduce participants' search times because they could move through the VE without worrying about collisions taking place. Second, with the stop algorithm, independent movement was expected to allow faster searching than the view-direction interface because participants could move directly in any direction to easily recover from collisions. A potential disadvantage of the slip algorithm was that it would increase the amount of VE sickness from which participants suffered because the algorithm could move them in a different direction to their indicated direction of travel.

3.1 Method

3.1.1 Participants. Twelve participants (three men and nine women) took part in the experiment, and their ages ranged from 18 to 24 years. All the participants were either undergraduates or graduates who volunteered for the experiment and were paid an honorarium for their participation. None of them had participated in experiment 1. Five participants withdrew and were replaced in the experiment. One simply did not wish to complete the experiment. The SSC data show that the others were suffering from nausea, but only two of them withdrew while using the slip algorithm.

3.1.2 Materials. The experiment used the same software and hardware as experiment 1, and, as before, participants moved around the VEs while physically standing up. The joystick was a small box (100×75×40 mm) that had a FASTRAK sensor mounted on the top and two buttons on the front. Participants used the joystick to move around the VE, but, as with the body-direction interface used in experiment 1, they could still look around without affecting their direction of

movement. To activate the joystick, participants pressed one of the buttons, and this caused the position of the joystick's center to be recorded (the origin). If the participants kept the button depressed and translated the joystick, then they moved along the vector that connected the origin to the joystick's new position, at a speed that increased with the distance between that position and the origin. The maximum speed (1.5 m/s) was achieved when the distance was 100 mm (beyond which the speed remained 1.5 m/s). Participants stopped whenever the button was released, and a new origin was defined when the button was pressed again. If the participants turned around while keeping the button depressed, then they traveled at a constant velocity relative to the heading recorded by the joystick's sensor (to achieve this, the VE software automatically redefined the joystick's origin), and this meant that, if the participants turned around at a constant rate, they would travel in a circle. Participants used the second button to select the targets. The view-direction interface worked in the same way as in experiment 1, but it used the joystick box for a device, rather than the pistol grip.

The slip algorithm was similar to that of Jacobson and Lewis (1997). When a collision occurred, the participants' desired movement was split into components that lay perpendicular and tangent to the colliding surface, and the algorithm moved them along the tangent component. (See figure 6.) Thus, participants would be unable to move only if they attempted to travel at exactly 90 deg. into a surface. This algorithm works well with convex obstacles, but, for environments that contain concave obstacles, a force-field algorithm, which works like repelling magnets, is more suitable. (See Xiao & Hubbold (1998).)

3.1.3 Procedure. The procedure was essentially the same as in experiment 1. Participants practiced each with interface and response algorithm before doing four test trials in each of the four conditions. As a precautionary measure, symptoms of VE sickness were monitored using the SSC.




Figure 6. Effect of the slip collision response algorithm when a person collides with a wall (top) and a cylindrical obstacle (bottom).

3.2 Results

As in experiment 1, the distributions of the distance, time, and path efficiency data were normalized using a logarithmic transformation. The collision data are not reported because there was little reason for participants to attempt to avoid the obstacles when the slip algorithm was used. Initial analyses showed no effect of trial number, so the data reported here were analyzed using two-factor (interface \times response algorithm), repeated-measures ANOVAs.




Figure 7. Mean time spent moving and stationary (excluding time spent selecting the targets) for each interface and response algorithm in experiment 2. Indep = independent movement interface. Error bars indicate standard error of the mean.

Participants traveled significantly less distance with the view-direction interface ($M = 52.8$ m, $SD = 12.3$) than with the independent interface ($M = 65.3$ m, $SD = 14.3$) ($F(1, 11) = 22.51$, $p < 0.01$). They also traveled significantly less distance with the slip algorithm ($M = 56.6$ m, $SD = 15.6$) than with the stop algorithm ($M = 60.4$ m, $SD = 15.9$) ($F(1, 11) = 22.51$, $p < 0.01$).

The time data are illustrated in figure 7. For the total trial time, participants were faster with the view-direction interface than the independent interface ($F(1, 11) = 11.89$, $p < 0.01$), and faster with the slip algorithm than with the stop algorithm ($F(1, 11) = 67.38$, $p < 0.01$). For the stationary time, there was no effect of interface ($F(1, 11) = 4.00$, $p > 0.05$), but participants were faster with the slip algorithm than with the stop algorithm ($F(1, 11) = 111.14$, $p < 0.01$). For movement time, participants were faster with the view-direction interface than with the independent interface ($F(1, 11) = 12.66$, $p < 0.01$), but there was no effect of response algorithm ($F(1, 11) = 0.02$, $p > 0.05$).

Analysis of the path efficiency data showed that participants revisited fewer cylinders with the view-direction interface than with the independent interface ($F(1, 11) = 10.61$, $p < 0.01$), and fewer cylinders with the




Figure 8. Mean percentage of possible target locations (the two-colored cylinders) that were revisited during each condition in experiment 2. Error bars indicate standard error of the mean.

slip algorithm than with the stop algorithm ($F(1, 11) = 30.27$, $p < 0.01$). (See figure 8.) Overall, participants revisited half (or more) of the cylinders in 21% of the trials and, as in experiment 1, all of the participants did this in at least one trial.

Two other sets of data are worth noting. First, participants used the full range of movements provided by the joystick, making 65% of their movement within ± 45 deg. of the forwards direction of their virtual body and 15% backwards (angle > 90 deg.). There was slightly more nonforward movement when the stop algorithm was used, compared with the slip algorithm, but the differences were small. Second, the slip algorithm reduced the amount of time that participants spent looking at their feet by a factor of 20. To see their virtual body (for example, in a collision), participants had to look downwards at an angle of at least 60 deg. This happened for 4% and 0.2% of the time when the stop and slip algorithms were used, respectively. Not only is this movement unnatural, evidence from our pilot studies and from the studies performed by others (J. R. Wilson, personal communication, February 21, 2000) suggests that changes of head pitch angle cause more symptoms of VE sickness than do changes of heading. Therefore, a slip algorithm may reduce levels of VE sickness even though the movement sometimes takes place in a different direction to which participants intended.

3.3 Discussion

Participants performed the trials quickest with the view-direction interface and the slip response algorithm. However, the primary reasons for the effects of interface and algorithm were different: the former was caused by a difference in the amount of time that participants spent moving, and the latter was caused by the amount of time that participants spent stationary. As well as taking less time, participants searched more efficiently with the view-direction interface and the slip algorithm. In other words, participants' performance was best when they used the simplest interface, despite the fact that it offered the least maneuverability. Contrary to expectations, independent movement (the virtual joystick) did not produce better performance than did view-direction movement when the stop algorithm was being used.

As in experiment 1, participants frequently retraced their steps. Comparison of the data for the view-stop condition in experiments 1 and 2 show that participants searched more slowly in the first experiment, but they also revisited fewer of the two-colored cylinders. If participants kept the speed button depressed during a collision, then their initial, post-collision speed would be nonzero, leading to faster movement but more collisions. Examination of the data showed that the mean value of this speed was much higher in the second experiment (0.49 m/s versus 0.27 m/s). It is likely that the speed-efficiency tradeoff was directly caused by the fact that participants used both stop and slip response algorithms in experiment 2, because, with the latter, there was little reason to release the speed button when a collision occurred.

The task used in the experiments would have been trivial to perform in the real world, yet, in the VEs, the ease of movement provided by the slip algorithm produced only a modest improvement in search efficiency. Clearly, another answer is needed to allow participants to move easily and effectively around cluttered VEs.

Part of this answer is likely to lie in participants' FOV, which should be substantially increased so they can see where they are in the environment as a whole and not just a keyhole view into the environment. Some real-world and VE studies that used small-scale spaces have

found effects of FOV on sketch map and layout reconstruction tasks (Alfano & Michel, 1990; Neale, 1997), although Alfano and Michel found a significant reduction in performance only when very narrow FOVs were used. The HMD used in experiments 1 and 2 had a specification that can be considered as fairly standard for current VE systems but, even so, its FOV (48×36 deg.) corresponded to a window of only 0.62 m by 0.45 m at arms' length (0.7 m). With a larger FOV, participants would be much less likely to travel past targets without seeing them, and would have to perform less view integration when learning a VE's spatial layout and memorizing where they had already traveled.

To investigate this, we performed a third experiment in which two separate groups of subjects navigated the narrow VE using either a normal (48 deg.) or wide (103 deg.) FOV. The width of the latter was chosen so that the vertical FOV was 90 deg., allowing participants to simultaneously see their body cylinder and the view that was horizontally ahead.

4 Experiment 3

Experiment 3 used a between-participants design: each participant used either the normal or wide FOV to search for targets in the narrow VE. All participants used the view-direction interface and slip collision response algorithm, the combination of which had produced the best performance in experiment 2. A desktop (monitor) display was used so that there was less chance of the conflict between the VE's geometric FOV and the display's physical FOV from causing symptoms of VE sickness.

4.1 Method

4.1.1 Participants. Twenty-four participants (twelve men and twelve women) took part in the experiment, and their ages ranged from 18 to 33 years. All the participants were either undergraduates or graduates who volunteered for the experiment and were paid an honorarium for their participation. None of them had participated in the other experiments, and they were

randomly allocated to the normal and wide FOV groups.

4.1.2 Materials. The experiment used the same software and hardware as experiment 1, but participants moved around the VE while physically seated in a chair. The view-direction interface worked in a similar way to experiments 1 and 2. Participants controlled their speed using one key on the keyboard and selected targets using another. Direction control was performed using the mouse. By using the mouse to move the cursor up and down the screen, participants could change the vertical view direction (pitch) by up to ± 90 deg. (zero-order control), and, by moving the cursor away from the center of the screen, the horizontal view direction changed by an amount that increased with the cursor's horizontal offset from the center (first-order control).

The monitor was viewed from approximately 60 cm, which meant that the physical FOV was similar to the graphical FOV used in the normal viewing condition. Unavoidably, there was distortion between the physical and graphical FOVs in the wide condition, but, if anything, this is likely to have reduced any advantage that was gained.

4.1.3 Procedure. The procedure was similar to the other experiments. Participants practiced the interface and then performed four test trials, using either the normal or wide FOV throughout.

4.2 Results

As in the other experiments, the distributions of the distance, time, and path efficiency data were normalized using a logarithmic transformation. Initial analyses showed no effect of trial number, so the data reported here were analyzed using between-participants ANOVAs.

Participants traveled a similar distance when using the normal ($M = 49.1$ m, $SD = 18.0$) and wide ($M = 42.9$ m, $SD = 10.1$) FOVs ($F(1, 22) = 1.49$, $p > 0.05$). For the time data (see figure 9), participants were significantly slower in total with the normal FOV than with the wide FOV ($F(1, 22) = 8.87$, $p < 0.01$), and




Figure 9. Mean time spent moving and stationary (excluding time spent selecting the targets) for each FOV in experiment 3. Error bars indicate standard error of the mean.

spent significantly more time stationary with the normal FOV than with the wide FOV ($F(1, 22) = 15.35$, $p < 0.01$). However, FOV had no effect on the time participants spent moving ($F(1, 22) = 3.47$, $p > 0.05$).

Analysis of the path efficiency data showed that participants revisited more two-colored cylinders with the narrow FOV than with the wide FOV ($M = 15.4\%$ versus 10.5%), but the difference was not significant ($F(1, 22) = 1.23$, $p > 0.05$). Inspection of the data showed that, in both conditions, participants completed approximately a quarter of the trials without revisiting any of the two-colored cylinders. However, trials in which a participant revisited half (or more) of the two-colored cylinders occurred three times more frequently with the normal FOV than with the wide FOV (19% and 6% of trials, respectively).

4.3 Discussion

In both conditions, participants frequently completed the trials by following an efficient path (little or no revisiting of the two-colored cylinders) but, with the wide FOV, there was a notable decrease in the proportion of trials in which participants appeared to get disoriented and substantially retraced their steps. Thus, as in the previous experiments, participants could complete the trials efficiently. However, a momentary lapse in

concentration was sufficient to cause disorientation, particularly with the normal FOV. A further performance advantage would be expected if the wide FOV was used with an HMD, because that type of display allows zero-order control for changes of view direction, which facilitates glances to one side, and a previous study has shown that use of an HMD causes participants to look around more while moving (Ruddle et al., 1999b).

There was a large difference in the amount of time that participants spent stationary in the two conditions. The cause of this time difference was the participants' need to perform much more view integration with the normal FOV. However, it is not known whether participants were performing this additional integration to plan where to travel or to maintain their memory for the places that they had already visited.

5 General Discussion

The primary aim of the present study was to investigate ways of allowing easy and efficient movement in cluttered VEs. The interfaces were parts of the VE system that were used by participants to achieve their goal (searching for targets), and this reflects the role of interfaces in VE applications. Of the three types of interface that were investigated (view-direction, body-direction, and independent movement), the simplest (view-direction) proved to be the most effective. This is contrary to the suggestions of Ruddle et al. (1997) and the findings of Bowman et al. (1997), which emphasized the (logical) importance of allowing VE users to look around while traveling, just as people do in the real world. In other words, decoupling the view and travel directions is only beneficial for some types of application. Independent movement was not found to be advantageous, but further investigations should be performed with other devices that allow this type of movement, and particularly those that allow zero-order control of position and orientation.

The slip collision response algorithm was effective and allowed participants to indicate where they wanted to travel without having to devote a large amount of effort, time, and attention to making small movements around

obstacles. Therefore, algorithms of this type should be implemented in cluttered VEs unless it is important to test the maneuvering ability of users.

A wide FOV is also important, and this has implications for the design of VE display systems. For example, few HMDs have a horizontal FOV that is wider than 50 deg. and many have one that is substantially narrower. Also, the keyhole view that is provided is one reason why HMDs are rarely used in commercial applications.

Finally, even when view-direction movement and the slip algorithm were combined with a wide FOV, participants sometimes became disoriented. Two parallels can be drawn between this occurring in the present study and the difficulties many participants have had in navigating large-scale VEs (Ruddle, in press). First, even when large-scale VEs contained many landmarks, some participants had such poorly developed memory for the paths they had traveled that they searched less efficiently than if they had chosen routes at random. Second, in both large-scale spaces and in the small-scale spaces used in the present study, participants had to integrate information that was seen in different views, over time. View integration is one factor that accounts for the large increase in the rate at which participants learned spatial knowledge from a map, compared with real-world navigation (Thorndyke & Hayes-Roth, 1982), and a similar effect was observed when small- and large-scale maps were used as navigational aids in VEs (Ruddle et al., 1999a). An understanding of the problem people have of searching small, cluttered VEs, and, in particular, the reasons why their memory for paths is so poor may lead to ways of significantly improving people's rate of spatial learning when they navigate large-scale VEs.

Acknowledgments

This work was supported by grant GR/L95496 from the Engineering and Physical Sciences Research Council. We thank Stewart Huddart for running the participants in experiments 1 and 2, and Justin Savage for running the participants in experiment 3. We gratefully acknowledge the anonymous reviewers for their helpful suggestions.

References

- Alfano, P. L., & Michel, G. F. (1990). Restricting the field of view: Perceptual and performance effects. *Perceptual and Motor Skills*, *70*, 35–45.
- Bowman, D. A., Koller, D., & Hodges, L. F. (1997). Travel in immersive virtual environments: An evaluation of view-point motion control techniques. *Proceedings of the Virtual Reality Annual International Symposium (VRAIS'97)*, 45–52.
- Card, S. K., Mackinlay, J. D., & Robertson, G. G. (1991). A morphological analysis of the design space of input devices. *ACM Transactions on Information Systems*, *9*, 99–122.
- Chance, S. S., Gaunet, F., Beall, A. C., & Loomis, J. M. (1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, *7*, 168–178.
- Cobb, S., V., G., Nichols, S., Ramsey, A., & Wilson, J. R. (1999). Virtual reality-induced symptoms and effects (VRISE). *Presence: Teleoperators and Virtual Environments*, *(8)*, 169–186.
- Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, *8*, 49–71.
- Gottschalk, S., Lin, M., & Manocha, D. (1996). OBB-Tree: A hierarchical structure for rapid interference detection. *Proceedings of the 1996 ACM Conference on Graphics (SIGGRAPH'96)*, 171–180. New York: ACM.
- Green, M. (1995). *The MR Toolkit Version 1.4* [Computer program]. Department of Computing Science, University of Alberta, Canada.
- Hand, C. (1997). A survey of 3D interaction techniques. *Computer Graphics Forum*, *16*, 269–281.
- Jacobson, J., & Lewis, M. (1997). An experimental comparison of three methods for collision handling in virtual environments. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*, 1273–1277.
- Neale, D. C. (1997). Factors influencing spatial awareness and orientation in desktop virtual environments. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*, pp. 1278–1282.
- Presson, C. C., & Montello, D. R. (1994). Updating after rotational and translational body movements: Coordinate structure of perspective space. *Perception*, *23*, 1447–1455.
- Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *15*, 1157–1165.
- Ruddle, R. A. (in press). Navigation: Am I really lost or virtually there? In D. Harris (Ed.), *Engineering psychology and cognitive ergonomics*. (Aldershot, UK: Ashgate)
- Ruddle, R. A., Payne, S. J., & Jones, D. M. (1997). Navigating buildings in “desk-top” virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, *3*, 143–159.
- . (1999a). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied*, *5*, 54–75.
- . (1999b). Navigating large-scale virtual environments: What differences occur between helmet-mounted and desktop displays? *Presence: Teleoperators and Virtual Environments*, *8*, 157–168.
- Templeman, J. N., Denbrook, P. S., & Sibert, L. E. (1999). Virtual locomotion: Walking in place through virtual environments. *Presence: Teleoperators and Virtual Environments*, *8*, 598–617.
- Thorndyke, P. W., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, *14*, 560–589.
- Weatherford, D. L. (1985). Representing and manipulating spatial information from different environments: Models to neighborhoods. In R. Cohen (Ed.), *The development of spatial cognition* (pp. 41–70). Hillsdale, NJ: Erlbaum.
- Witmer, B. G., Bailey, J. H., Knerr, B. W., & Parsons, K. C. (1996). Virtual spaces and real-world places: Transfer of route knowledge. *International Journal of Human-Computer Studies*, *45*, 413–428.
- Xiao, D., & Hubbard, R. (1998). Navigation guided by artificial force fields. *Proceedings of the Computer Human Interfaces Conference (CHI'98)*, 179–186.