Published article:

Ruddle, RA, Howes, A, Payne, SJ and Jones, DM (2000) *Effects of hyperlinks on navigation in virtual environments.* International Journal of Human Computer Studies, 53 (4). 551 - 581.

http://dx.doi.org/10.1006/ijhc.2000.0402

The Effects of Hyperlinks on Navigation in Virtual Environments

Roy A. Ruddle

Andrew Howes

Stephen J. Payne

Dylan M. Jones

School of Psychology, Cardiff University

PO Box 901, Cardiff, UK. CF11 3YG

{ruddle, howesa, paynes, jonesdm}@cardiff.ac.uk

Corresponding author: Roy Ruddle [contact author for other figures]

Ruddle, R. A., Howes, A., Payne, S. J., & Jones, D. M. (2000) The effects of hyperlinks on navigation in virtual environments. *International Journal of Human Computer Studies, 53*, 551-581.

Summary

Hyperlinks introduce discontinuities of movement to 3-D virtual environments (VEs). Ten independent attributes of hyperlinks are defined and their likely effects on navigation in VEs are discussed. Four experiments are described in which participants repeatedly navigated VEs that were either conventional (i.e., obeyed the laws of Euclidean space), or contained hyperlinks. Participants learned spatial knowledge slowly in both types of environment, echoing the findings of previous studies that used conventional VEs. The detrimental effects on participants' spatial knowledge of using hyperlinks for movement were reduced when a time-delay was introduced, but participants still developed less accurate knowledge than they did in the conventional VEs. Visual continuity had a greater influence on participants' rate of learning than continuity of movement, and participants were able to exploit hyperlinks that connected together disparate regions of the a VE to reduce travel time.

1. Introduction

When people travel through a conventional virtual environment (VE; a virtual reality world) they experience a continuous 3-D space that has been constructed using real-world (Euclidean) laws of spatial structure. Hyperlinks allow discontinuous movement (direct movement between places that are far apart) and, therefore, break some of those Euclidean laws but confer potential advantages over conventional VEs in terms of the design of flexible layouts and reduced navigation time. These are the same advantages that hypermedia have over conventional books. The disadvantage of hyperlinks in VEs is that they are likely to add the cognitive difficulties that are known to be caused to people by discontinuities in hypermedia (Conklin, 1987; Kim & Hirtle, 1995; Utting & Yankelovich, 1989) to the already considerable problems of navigating conventional VEs (e.g., see Ruddle, Payne, & Jones, 1997).

Implementations of hyperlinks in VEs are becoming increasingly common and include the Uniform Resource Locators that connect together groups of VEs which are stored on the World-Wide Web, for example, Superscape's Virtual World-Wide Web (http://vwww.com), and all multi-node movies that are constructed using formats such as Apple Computer's QuickTime VR. Despite this, there is almost a complete lack of published experimental research that has studied the effects of hyperlinks, in their various forms, on navigation in VEs as opposed to text-based hypermedia.

In this article we identify 10 attributes of hyperlinks and outline their likely effects on navigation in VEs. Four experiments are described which investigated some of these attributes. Taken together, the attributes and the experiments provide a baseline for studies of navigation in this type of VE. First, however, we set the scene by summarising studies that have investigated navigation in conventional VEs.

2. Navigation in Conventional VEs

In recent years, many experimental studies have investigated aspects of navigation and the development of spatial knowledge in conventional, large-scale VEs (large-scale environments are those in which a person cannot resolve all the detail necessary for efficient navigation from a single human's-eye viewpoint, see Weatherford, 1985). Those studies have shown that the navigation of conventional VEs relies on and is associated with cognitive processes that are similar to those that people employ in real-world settings. For example, people form similar spatial representations when they navigate VEs and real-world environments (Hancock, Hendrix & Arthur, 1997; May, Péruch & Savoyant, 1995; Tlauka & Wilson, 1996). People ultimately learn to wayfind efficiently (e.g., travel by the shortest route) and develop accurate survey (map-type) knowledge (Ruddle et al., 1997), and can successfully transfer the knowledge that they learn in a VE to an equivalent real-world environment (Bliss, Tidwell & Guest, 1997; Witmer, Bailey, Knerr & Parsons, 1996). In other words, people are not always disoriented when they navigate VEs, as was sometimes thought, but they do often take a long time to develop accurate spatial knowledge.

A variety of techniques for reducing that time have been investigated. One of these, supplementary landmarks, causes people change their navigational strategy (Darken & Sibert, 1993, 1996; Tlauka & Wilson, 1994) but experimental data suggest that this sometimes produces only a modest improvement in the accuracy of people's way-finding and no improvement in their survey-type knowledge (Ruddle et al., 1997).

Another technique is to provide an interface that allows normal physical movements to be made (e.g., head-tracking) rather than the abstract interface provided by a mouse and keyboard. When people use an HMD some aspects of their spatial knowledge develop more quickly, but there is no difference in the accuracy of their wayfinding (Ruddle, Payne, & Jones, 1999b). However, larger differences seem to occur if an interface allows people to actually walk around a VE (physical translational and rotational movements; Chance, Gaunet, Beall, & Loomis, 1998; Grant & Magee, 1998).

As has already been stated, conventional VEs are constructed using real-world rules of spatial structure, and this can be a disadvantage for some applications. Hyperlinks can reduce the distance and travel time between disparate regions of a VE, and make different places appear to be spatially related, even if they actually exist at completely separate structural locations (cf. the lists of “other useful links” that are found on some WWW pages). The effect that hyperlinks have on a person’s ability to navigate, or the rate at which they learn a VE’s spatial structure, is likely to be affected by the hyperlinks’ attributes and these are discussed in the following sections.

3. Hyperlinks in VEs

In a conventional VE a person travels in small discrete steps at a time interval that is equal to the VE’s frame rate. For example, a person who travels at a normal walking speed (5 km/h) through a VE that has a frame rate of 20 Hz., travels approximately 20 mm between frames. Given the bounds of the human perceptual system, this effectively means that the person experiences each and every place along the paths that they travel, which helps to ensure that they perceive their movement as continuous. We refer to this property as experiencing ∂-place.

Hyperlinks allow a person to jump from one location to another without experiencing ∂-place. These devices are similar in concept to the links that occur within hypermedia. Therefore, hyperlinks are likely to compound the substantial difficulties that a person has navigating conventional VEs with the problems that occur when a person navigates hypermedia, for example, disorientation and cognitive overhead (respectively, the tendency of a person to lose their sense of location and direction, and the additional effort and concentration that is necessary to maintain several trails at one time; Conklin, 1987). However, set against those potential disadvantages are the advantages of speed and design flexibility that were outlined above.

By way of illustration, a real world analogy to the perceptual consequences of hyperlinks in VEs can be found in underground (subway) railway networks. When a person travels on an underground railway their experience of ∂-place is limited to views of dark tunnel walls that give little impression of where or how far they have travelled. When they arrive at a new station they need other information, for example, prior knowledge, or a map, to know their spatial location relative to the station from which they started.

In general, the severity of the navigational problems and the size of any advantages are affected by the characteristics of the hyperlinks. In the following sections, these characteristics are defined using 10 independent, attributes (see **Table 1**).

3.1. Integration

Integration refers to the way in which hyperlinks are combined with other types of movement (see **Figure 1**). Where hybrid integration is used, hyperlinks supplement movement through VEs that in all other ways are conventional. Every node has a known, unique position in Euclidean space because a person may travel to it by moving continuously. Hyperlinks are likely to be integrated in this way to facilitate rapid movement between nodes that are a large straight-line distance apart, or to provide instantaneous access to important nodes.

Hyperlinks may also connect regions of conventional virtual space that otherwise are separate. These environments can be thought of as consisting of a number of Euclidean islands that are joined by hyperlink bridges.

A person may structure their knowledge of hybrid VEs in similar ways to conventional VEs and may, if they wish, explore these VEs without using the hyperlinks. VEs that contain bridges may only be fully explored if the bridges are traversed. This forces a person to experience the discontinuities of movement (something that is likely to make navigation more difficult) but divides the VEs into smaller and perhaps more manageable islands of space (distinct districts have been shown to be a desirable feature for cities that are easy to navigate; Lynch, 1960).

3.2. Appearance

Hyperlinks may have the appearance of physical objects, or they may be implemented using other devices such as a link menu, or a hot-list. An object is displayed within a VE at a particular spatial position and, therefore, must be located before the associated hyperlink may be traversed. A link menu is a “non-physical” device that is associated with a localised region in a VE and provides access to hyperlinks whenever a person is in the link menu’s region. A hot-list is also a non-physical device but, effectively, it is external to the VE’s structure and, therefore, may be accessed from any position.

3.3. Time

The time that is required to traverse hyperlinks may be implemented in a number of different ways. The first of these is to minimise the time so that a person moves from node to node in one frame (instantaneous movement). A second implementation is to set the traversal time equal to the time that would be required to travel the Euclidean straight-line distance between the ends of a hyperlink, and that "continuity" of time might help a person to form an accurate mental model of the VE. On the other hand, that increase in traversal time would introduce time delays which might reduce a person's ability to spatially connect sequences of nodes and lead to performance reductions that are similar to those which occur when a person performs tasks in VEs which operate at slow frame rates (e.g., Arthur, Booth & Ware, 1993; Richard, Birebent, Coiffet & Langrana , 1996).

A third type of implementation is to use a fixed (non-instantaneous) time. An example of this is the links which connect VEs that are accessed via the WWW, although it should be noted that the link traversal time also depends on factors such as the size of the files which are to be retrieved and the amount of network traffic. O'Hara and Payne (1998) found that participants performed more planning and, therefore, searched more efficiently when time-delays were introduced into the user interface to various puzzles. A similar effect might be found by introducing delays to the amount of time required to traverse hyperlinks.

3.4. Motion profile

Hyperlinks which allow instantaneous traversal give a person no perception of motion (they jump between nodes). Although this provides the fastest mechanism of traversing hyperlinks, other motion profiles should be considered. In particular, a profile that moves a person at a small velocity for several frames at the start and end of jumps would provide some optical flow and, therefore, information about the resultant (or implied) direction of travel. Hyperlinks in QuickTime movies use an implementation of variable velocity motion profile.

3.5. Preview

When a person travels through conventional VEs they can sometimes see either their desired destination from far away, signposts which indicate routes to other, related nodes, or visual information that acts as landmarks and aids efficient wayfinding (c.f., residue; Furnas, 1997). A hyperlink may provide preview by indicating the node to which it leads prior to traversal (single-step preview). This provides localised information that increases the area within which a node may be detected. Single-step preview allows the content of nodes to be determined before hyperlinks are traversed and the relative locations of two or more nodes to be directly determined from a single place. To minimise clutter, signposts are most likely to indicate the routes to a small number of important nodes. However, this world-based information may help a person to develop knowledge about the overall structure of an environment, as well as particular routes.

3.6. Visual Continuity

Visual continuity (also called visual momentum; Wickens, 1992; Woods, 1984) refers to the quantity and effectiveness of perceptual anchors which are present in the images that are displayed when a person traverses hyperlinks, and which help a person to maintain knowledge of their position and orientation. In conventional VEs the visual continuity between images that are rendered on a display is nearly 100%, but if the nodes at each end of a hyperlink are in completely different parts of a VE then the visual continuity is zero (nothing appears in both images). Other hyperlinks may connect nodes that either lie within sight of each other, or that both lie within the line of sight of the same global (external) landmark, for example, a distant building (Evans, Skorpanich, Gärling, Bryant & Bresolin, 1984). In these latter situations some visual continuity is provided by cues which are present in the images that are displayed at both ends of a hyperlink, and this helps a person to perceive their movement. The resultant navigational benefits will depend on the amount and the prominence of the continuity between the images.

3.7. Directionality

Links may be unidirectional or bidirectional. Unidirectional links are likely to lead to the type of navigational problems that a person encounters when they drive vehicles along one-way road systems (a person may be able to see their destination but have difficulty working out how to reach it) or use interfaces that do not contain an 'undo' option. However, designers could exploit the characteristics of unidirectional links in conventional or discontinuous VEs to guide (or force) a person to travel along certain routes and, therefore, to pass specific features.

3.8. Directional Affordance

Directional affordance refers to differences between a person's implied and actual direction of travel when they traverse hyperlinks. In **Figure 2** the hyperlinks that are situated in Room A are positioned against the walls. A person who stands in this room might assume that the hyperlinks lead to rooms that lie in the same direction as the hyperlinks (after all, this is what happens when a person walk through doors in conventional VEs and in the real world). When a person traverse hyperlink H_c they actually travel in the opposite direction. In VEs such as multi-story buildings, hyperlinks that imply sideways movement could, in fact, move a person up or down. Hyperlinks that are selected using non-spatial devices such as link menus and hot-lists are not displayed spatially and, therefore, have indeterminate direction.

One way of quantifying directional affordance is by using vector algebra. If the dot (scalar) product of the unit vector from a person's current position to a hyperlink and the unit vector from the person's current position to their next (i.e., post-hyperlink) position is 1.0 then the directional affordance is 100% consistent (hyperlink H_b). By contrast the dot product for hyperlink H_c is -1.0.

3.9. Rotational Offset

The size of a rotational offset is the angular change in a person's direction of view that happens when they traverse a hyperlink, but which they do not experience in terms of gradual changes in the scene that is displayed. In **Figure 3**, the rotational offset of hyperlink H_0 is 0° because the person's direction of view remains unchanged when they traverse it, but hyperlink H_{90} has a directional offset of 90°. In effect, the orientation of the person's frame of reference changes when they traverse H_{90}.Although links that contain rotational offsets are very likely to lead to disorientation and to make it difficult for a person to navigate, they do occur in some commercial applications (e.g., Virtual Tourism: Paris, 1996). The maximum possible angle of a rotational offset in a 2-D plane is 180°.

3.10. Crossings

Non-planar VEs are those in which some hyperlinks intersect but do not connect, whereas planar VEs contain none of these crossings. Crossings produce the "spider's web" of connections that is a characteristic of hypermedia and produce visual clutter on overview maps, but can be used to increase design flexibility ("adjacent" places no longer have to fit within the confines of Euclidean space) and reduce travel time, particularly between nodes that are a large Euclidean distance apart.

The effect that hyperlinks have on the distance between the nodes of a VE depends on the VE's structure. For example, if a 100-node environment is arranged in a line (a spatial structure that is easy to describe, visualise, and navigate) then the maximum distance between nodes is 99 units. The addition of a single, optimally-placed additional hyperlink reduces this to 50 units. If, instead, that environment is arranged as a grid (the 2D layout that is used for the streets in places such as Manhattan and Milton Keynes), then a single, optimally-placed additional hyperlink reduces the maximum distance between nodes from 18 to 14 units. For a 5 X 5 X 5 cube, an individual hyperlink can only reduce the maximum distance between nodes from 12 to 11 units. If sufficient hyperlinks are added a VE becomes fully-connected (each node is directly connected to every other node) and the mean and maximum distances become 1 unit. Of course this produces a network of connectivity which is unmanageable in all but the smallest VEs.

The remainder of this article describes four experiments that investigated the effects of time, visual continuity, and crossings in hyperlinks on participants' navigation in virtual buildings. In the first three experiments the hyperlinks connected together rooms that were adjacent in Euclidean space, and this allowed the effects of hyperlinks to be studied in isolation of environmental features such as crossings. Crossings were implemented in the final experiment to study the effects of "useful" hyperlinks, which substantially reduced the time required to travel between certain rooms. Videos that show examples of the door and hyperlink connections used in the experiments are stored in QuickTime VR format at http://www.cf.ac.uk/uwc/psych/ruddle/C-HIVE/DVE/.

The first experiment used a simple form of hyperlink which, when traversed, made an almost-instantaneous change in a person's position. The views of the VE at the two ends contained nothing in common, which meant there was zero visual continuity. The impact this can have was demonstrated by Bowman, Koller, and Hodges (1997). In their study, participants took longer to re-orient themselves when they jumped between positions in a VE than when they traveled continuously, but the effect this movement had on navigation was not investigated.

4. Experiment 1

For Experiment 1 two versions of three different virtual buildings (Buildings A, B1 and B2; see **Figures 4** and **5**) were created. Building A was used for a practice test, and Buildings B1 and B2 were used for the full experimental tests. In each building, the connectivity of the rooms was the same in both versions, but in one the connections were doors and in the other the connections were hyperlinks. The hyperlinks were an 'H' shape that was the same size as the doors (see **Figure 6**). Participants repeatedly navigated one version of each building.

4.1. Method

4.1.1. Participants

Twenty-four participants (15 men and 9 women) took part in the experiment. All the participants were either students or graduates, who volunteered for the experiment and were paid an honorarium for their participation. Their ages ranged from 19 to 27 years (M = 22.1). Each participant first navigated Building A. Then they either navigated Building B1, followed by Building B2, or they navigated Building B2 and then Building B1. Each participant was randomly assigned to one of eight groups, which were used to counterbalance the type of connection used in Building A, and the order of navigation and type of connection used in Buildings B1 and B2. Two participants asked to withdraw, one in the first building and the other in the second building, and were replaced in the experiment.

4.1.2. Materials

The experiment was performed on a Silicon Graphics Crimson Reality Engine, running a C++ Performer application that we designed and programmed. A 21-in. (53 cm) monitor was used as a display and the application update rate was 20 Hz.

The buildings contained chessboard-type arrangements of rooms (see Figures

4 and 5). This type of layout was chosen so that the effects of the different types of connection could be studied in isolation of environmental factors typically present in VEs that model real buildings, for example, orientation and position information that is provided by views of distant objects down long corridors. Pilot studies were used to help determine an appropriate number of rooms for the VEs, and hence one component of navigational complexity. Barriers to movement are another component of complexity, and these were created by omitting the connections between some adjacent rooms. To travel between each object pair used in the fourth stage of the test procedure (see below) a participant had to enter a minimum of three rooms (excluding the starting room). The connections that were omitted were chosen so that each pair was connected by a unique, shortest route.

Each door in the VEs was opened automatically by an amount that depended on a participant's distance from the door. The door was closed if the participant was more than 2 m from the door and fully open if the participant stood in the doorway. All the hyperlinks were two-way (participants could retrace their steps). Each room of each building contained a scaled, 3-D model of a different everyday object and these were placed on top of gray box that was 1 m in height (see Figure 6). Each object was made visible when a participant entered the object's room, and invisible when the participant left the room. This was done so that the participant could not look through a partly open door to see which object was in the adjacent room, something that could not be done anyway if the rooms were connected by hyperlinks.

To define what was seen on the monitor, the application had to specify the field of view to be used and the height above the buildings' floor at which viewing took place (effectively a participant's virtual "eye" height). The FOV was set to be 90° and each participant's virtual eye height was set equal to their actual eye height.

The center of the screen was marked using a small green square.

The interface used the mouse and five keys on the keyboard, and allowed a participant to look around while traveling in a straight line. If the participant held down the left or right mouse buttons, then the azimuth of their view direction changed at a rate of 60 degrees/s. By moving the mouse to change the offset of the cursor from the center of the screen, the participant could vary the azimuth and altitude of their direction of view by ± 90°. Four of the keys allowed the participant to slow down, stop, speed up, and move at the maximum allowed speed (4.8 km/h). The fifth changed the participant's direction of movement to the current view direction. Participants were prevented from traveling through walls by a collision detection algorithm.

To traverse a hyperlink a participant placed the cursor anywhere on the 'H' and then pressed the middle mouse button. The same mechanism was used to select the target objects (see below).

4.1.3. Procedure

Participants were run individually. They were told that the experiment was being performed to assess people's spatial knowledge in VEs, all the VEs had one story, some of the VEs might contain discontinuities, and that their primary goal in the tests was to enter as few rooms as possible. Rooms were counted each time they were entered. In the experiment, a participant first practiced the interface controls, and then navigated the three test VEs. The first of these (the practice test; Building A) familiarized participants with the procedure used in the test and reduced the general learning effect that was expected to be present between the two full tests (see Ruddle, Payne & Jones, 1998; Stanton, Wilson & Foreman, 1996). The elapsed time for each participant's practice and tests was approximately 3 hr.

Interface practice. Two buildings were used. In one the rooms were connected by doors and in the other the rooms were connected by hyperlinks. Both buildings contained five rooms but their layouts were different. Participants used the first and second practice buildings respectively to familiarize themselves with the controls for moving around rooms and though doors, and traversing the hyperlinks.

Test procedure. The test procedure was the same for all three test buildings and was divided into five stages. A participant completed all five stages in one building before resting for 3 min and then starting the test in the next building. For Stage 1, a participant was taken into a one-room virtual building which contained a number of objects. These were the objects that were designated as targets (6 objects in Building A, and 8 objects each in Buildings B1 and B2). When the participant had familiarized themselves with the appearance of the objects they started the next stage.

Stages 2 and 3 were identical. All participants started in the same room (see Figures 4 and 5) and searched for all the target objects in any order. These searches were designed to investigate how efficiently a participant could search the whole of a VE (the first search was uninformed, because the participant had no prior knowledge of the VE's layout, and the second was informed; this terminology is consistent with Ruddle, Payne, & Jones, 1999a). At all times a message on the monitor showed the name of the object(s) that were still to be found. When an object was found and selected its name was deleted from the screen.

In Stage 4 the participant had to repeatedly find specific target objects in a test that was designed to investigate whether the participant could learn a specific route (a repeated-route) through the VE and take short-cuts. The objects were divided into pairs, with one of each pair designated as the start and the other as the finish (see Figures 4 and 5). For each group of participants one pair was used to test the learning of the repeated route and the other pairs were used to test the short-cuts. Different groups used a different pair for the repeated route. The detail of the procedure was quite complex, and is best illustrated by using Building B1 and participants in Group 1 as an example.

A participant in this group started in the center of the room that contained the saucepan. A message on the screen said “Revisit target saucepan (it is beneath you)” and the participant used the mouse to select the saucepan. This caused the message to change to “Revisit target helicopter”. The participant then searched the VE until they found and selected the helicopter. This caused the screen to go blank for 5 s, and when the VE was displayed again the participant was in the room that contained the house and the message said “Revisit target house (it is beneath you)”, and so on. In total, the participant searched for eight pairs of objects in the following order: saucepan/helicopter, house/traffic light, saucepan/helicopter, clock/nail, saucepan/helicopter, pencil/axe, saucepan/helicopter, house/traffic light, saucepan/helicopter, clock/nail, saucepan/helicopter, and pencil/axe. The saucepan/helicopter pair tested the repeated-route learning and the other pairs tested the participant’s ability to take short-cuts. Although the terms repeated-route and short-cut are used to discriminate between the two types of pair, the real difference was the frequency with which, and number of times that, participants performed the searches.

For the final stage (Stage 5) the participant was asked to draw a sketch map of the building. They were given written instructions which asked them to draw as much detail as possible, including the names of the objects, and the doors and hyperlinks that connected the rooms.

4.2. Results

Only the data for the full tests (Buildings B1 and B2) will be reported. The any-order, repeated-route and short-cut searches were measured by recording the number of rooms that participants entered and the time that they took. The data for the three short-cut pairs were averaged for each search number. The distribution of the number of rooms data was normalized using a square root transformation, and the time data was normalized using a logarithmic transformation. Each participant navigated either Building B1 or Building B2 with doors, and the other building with hyperlinks. For the analyses reported below these data are collectively termed as referring to Building B. The data were analyzed using mixed analyses of variance (ANOVAs) that treated the search number and the Building B connection type as repeated measures, and the type of connection participants had used in the practice test (Building A) as a between-groups factor. Interactions and the effect of the Building A connection type are only reported if they were significant.

4.2.1. Number of Rooms Entered

Participants entered significantly fewer rooms when they navigated using doors than when they used hyperlinks in the any-order searches, $F(1, 22) = 8.50$, $p = .01$, the repeated-routes, $F(1, 22) = 4.90$, $p = .04$, and the short-cuts, $F(1, 22) = 8.57$, $p = .01$ (see **Table 2** and **Figure 7**). The difference between the informed and uninformed any-order searches was not significant, $F(1, 22) = 3.26$, $p = .08$, but there were main effects of search number for the repeated-routes, $F(5, 22) = 5.04$, $p < .01$, and the short-cuts, $F(1, 22) = 4.77$, $p = .04$.

4.2.2. Time Taken

In general, the time and number of rooms data showed the opposite effects to each other. Participants took more time when they navigated using doors than when they used hyperlinks in the any-order searches, $F(1, 22) = 50.06$, $p < .01$, the repeated-routes, $F(1, 22) = 19.47$, $p < .01$, and the short-cuts, $F(1, 22) = 17.54$, $p < .01$ (see **Table 3** and **Figure 8**).

There were main effects of search number in the any-order search, $F(1, 22) = 45.36$, $p < .01$, the repeated-routes, $F(5, 22) = 6.31$, $p < .01$, and the short-cuts, $F(1, 22) = 10.46$, $p < .01$. Also, participants who had navigated the doors version of Building A navigated both versions of Building B significantly more quickly in all three searches: any-order (M = 162 s vs. 273 s), $F(1, 22) = 16.21$, $p < .01$, repeated-route (M = 37 s vs. 50 s), $F(1, 22) = 4.23$, $p = .05$, and short-cuts (M = 51s vs. 65 s), $F(1, 22) = 6.77$, $p = .02$.

4.2.3. Random Walk Analysis

Participants sometimes had great difficulty finding the target objects, particularly in the any-order searches. To further explore the efficiency of their searches we compared the number of rooms that participants entered with the number of rooms that would have been entered by a hypothetical participant who chose at random which rooms to enter. A computer program was written to generate 1,000,000 random walks for Building B (Buildings B1 and B2 were mirror images of each other and, therefore, could be combined into one analysis) and thus determine the percentage of "random" participants who would have completed the any-order search by entering a given number of rooms. The number of rooms that participants in the experiment entered in Building B was sorted into an ascending order, with the sorting performed separately for the two types of connection. These data, illustrated in **Figure 9**, show that, in the doors condition, participants searched more efficiently than a random walk, but in the hyperlinks condition participants searched less efficiently than the random walk.

4.2.4. Sketch Maps

The sketch map data were designed to supplement the objective navigation data (number of rooms and time taken), and to investigate whether fundamental differences occurred in the structure of participants' spatial representations when they navigated VEs that contained either doors or hyperlinks.

Five metrics were used to calculate scores from each sketch map. The first metric (object names) was the number of objects that a participant correctly named or drew (some participants seemed to be talented artists!). The second (cells) measured participants' knowledge of the spatial structure of a VE, independent of the objects. The VE's actual layout was normalized for scale, overlaid on a sketch map and rotated to maximize the number of rooms that were correctly positioned relative to each other. The score was this maximum number of rooms, minus the number of rooms that were incorrectly positioned. The third metric (spatial names) used a similar method of measurement to the cell score metric, but took into account the name of the object in each room. Unnamed and incorrectly named rooms were ignored. The fourth metric (adjacent rooms) counted the number of named rooms that were correctly drawn in spatially-adjacent positions to each other, a maximum of 24. The fifth metric (links) counted the number of links (i.e., doors or hyperlinks) between named rooms that were correctly drawn, a maximum of 19. Once the maps had been analyzed each participant's scores were converted to give the percentage that were correct for each metric.

Most of the sketch maps that participants drew showed the rooms in spatially adjacent locations, usually as a grid of squares. Participants mean scores in the doors and hyperlinks conditions were: 82% vs. 85% (object names), 82% vs. 73% (cells), 46% vs. 35% (spatial names), 43% vs. 37% (adjacent rooms), and 40% vs. 38%

(links). The data were analyzed using between participants ANOVAs. A significant effect of connection type occurred for the cells, $F(1, 23) = 4.36$, $p = .05$, and a marginally significant effect occurred for the spatial names, $F(1, 23) = 3.48$, $p = .07$.

4.3. Discussion

In Experiment 1 participants' goal was to enter as few rooms as possible. By this metric, participants were more efficient in all of their searches when they used doors than when they used hyperlinks. The provision of hyperlinks to connect rooms in the VEs seemed to inhibit the rate at which participants developed spatial knowledge. The sketch map data, and in particular data from the cells metric, support this finding and indicate that participants developed more accurate knowledge of a VE's overall layout when doors were used as the connections.

The primary differences between the two types of connection were: (a) Participants moved continuously with the doors but jumped from room to room with the hyperlinks (in the latter they did not experience δ-place), (b) Movement in the doors condition was visually continuous but there was zero visual continuity each time a participant traversed a hyperlink, and (c) It took approximately 4 s to "walk" across a room but hyperlinks were traversed almost instantaneously (one graphics frame, 0.05 s). It is possible that all of these factors contribute to the differences in search efficiency and spatial knowledge that were observed.

The time data highlighted the effect of the different door and hyperlink traversal times. Participants performed all of the searches in Building B more slowly when they used doors than when they used hyperlinks. In other words, their searches through the hyperlinks version were more efficient in terms of time, despite being less efficient in terms of distance (number of rooms). The time taken to traverse a room, and hence the cost of a navigational error, was greater in the doors condition than in the hyperlinks condition. Despite being instructed to try and minimize the number of rooms that were entered, the data suggest that participants used different strategies in the two conditions, and with hyperlinks their strategy seemed more akin to a random walk than purposeful planning. In fact, as Figure 9 shows, in their initial searches participants tended to perform slightly worse in the hyperlinks condition than they would have if they had chosen hyperlinks at random.

Within the data for Building B, a consistent time advantage was exhibited by the participants who had navigated the doors version of Building A. A likely explanation is that those participants gained a better understanding of the type of building layouts being used in the study, and that gave them a spatial framework on which to base their navigation.

In summary, one reason for the differences that occurred between navigation with doors and hyperlinks was probably the change in participants' strategy that was caused by the difference in the time penalty of navigational errors. This is consistent with the findings of O'Hara and Payne (1998). This strategy change is important for practical implementations of hyperlinks because most commercial applications are likely to minimize the traversal time (after all, nobody likes being held up). However, to compare the effects of doors and hyperlinks in isolation of the time cost of errors we performed a second experiment in which each hyperlink took an extended amount of time to traverse.

5. Experiment 2

5.1. Method

5.1.1. Participants

Twenty-four participants (15 men and 9 women) took part in the experiment. All the participants were either students or graduates, volunteered for the experiment, were paid an honorarium for their participation, and had not taken part in Experiment 1. Their ages ranged from 17 to 27 years ($\underline{M}$ = 22.7). As in Experiment 1, each participant was randomly assigned to one of eight groups. One participant asked to withdraw after drawing the sketch map for Building A and was replaced in the experiment.

5.1.2. Materials and Procedure

Experiment 2 used the same virtual buildings, software and procedure as Experiment 1. The only difference between the two experiments was what happened when a participant traversed a hyperlink. In Experiment 1 the traversal took one graphics frame (0.05 s), but in Experiment 2 the traversal took 4 s (equal to the amount of time that it took to walk across one of the rooms). When a participant in Experiment 2 selected a hyperlink a message was displayed on the screen and which said "Please wait" and the screen froze. After 4 s the message was removed and the participant was jumped to the room that was at the other end of the hyperlink.

5.2. Results

As in Experiment 1, only the data for the full tests were analyzed. There was no main effect of connection type in the any-order searches, $\underline{F}(1, 22) = 1.85$, $\underline{p} = .19$, but participants did improve from the uninformed to the informed search, $\underline{F}(1, 22) = 7.55$, $\underline{p} = .01$ (see **Table 4**). Planned contrasts showed that the difference between doors and hyperlinks was significant for the uninformed search, $\underline{F}(1, 22) = 5.26$, $\underline{p} = .03$, but not for the informed search, $\underline{F}(1, 22) = 0.03$, $\underline{p} = .85$.

For the repeated-route searches participants entered a similar number of rooms with both types of connection (see **Figure 10**). There was no main effect of connection type, $\underline{F}(1, 22)$ = zero, $\underline{p} = .95$, but there were significant effects of search number, $\underline{F}(5, 22) = 12.80$, $\underline{p} < .01$, and the type of connection that participants had used in Building A (doors: M = 4.7; hyperlinks: M = 6.8), $F(1, 22) = 7.53$, $p = .01$.

For the short-cut searches there was no effect of connection type, $F(1, 22) = 1.00$, $p = .33$, but, as with the repeated-route searches, there was a significant effect for the type of connection that participants had used in Building A (doors: M = 6.8; hyperlinks: M = 10.1), $F(1, 22) = 5.74$, $p = .03$.

5.2.1. Comparisons with Experiment 1

To further investigate the effects of introducing a time-delay to the hyperlinks we performed additional analyses which compared the number of rooms that participants entered in the hyperlinks condition of Experiments 1 and 2. The analyses were performed using repeated measures ANOVAs, with the type of hyperlink treated as an independent variable. In both experiments the type of connection (doors or hyperlinks) that participants used in Building A affected their performance in Building B. To remove that effect the analyses were performed using only data for the 50% of participants who used the same type of connection in Building A and the first of the two full tests. In all three types of search participants in the time-delay condition entered fewer rooms than participants in the instant condition. The difference was significant for the repeated-routes (M = 7.6 vs. 10.5), $F(1, 10) = 5.56$, $p = .04$, but not for the short-cuts (M = 9.7 vs. 11.0), $F(1, 10) = 0.30$, $p = .60$, or the any-order searches (M = 36.3 vs. 40.6).

5.2.2. Sketch Maps

The sketch map data were analyzed using the same five metrics and ANOVAs as Experiment 1. Participants mean scores in the doors and hyperlinks conditions were: 84% vs. 90% (object names), 77% vs. 72% (cells), 56% vs. 60% (spatial names), 52% vs. 61% (adjacent rooms), and 50% vs. 50% (links). The differences were marginally significant for the object names, $F(1, 23) = 3.75$, $p = .07$, and the adjacent rooms, F(1,

23) = 3.72, $p = .07$.

5.3. Discussion

The introduction of a time delay to the hyperlinks reduced, but did not eliminate, the differences between navigation in a virtual building that was connected using doors and navigation using hyperlinks. In general, the differences between the doors and hyperlinks conditions were most significant during the early stages of searching, and participants quickly adapted to navigating using hyperlinks. One contributory factor may have been the structural similarities between the VEs.

An interesting turn around took place in the sketch map data. In Experiment 1, participants' sketch maps tended to be more accurate in the doors condition than in the hyperlinks condition. In Experiment 2 the opposite was true, although the differences were not statistically significant, and it seems that participants in the hyperlinks condition may have been learning the layout of the VEs more in terms of adjacent rooms than the VEs' overall spatial form.

As well as the time required to traverse the different types of connection, two other probable causes of the differences between doors and hyperlinks were identified in the discussion section of Experiment 1, namely the effects of movement continuity and visual continuity. To further investigate these we performed a third experiment in which the doors in the VEs did not open but could be walked through (in effect, the participants acted as if they were ghosts). These ghost-like doors looked identical to and preserved the movement continuity of the doors used in Experiments 1 and 2 but, like the hyperlinks, had zero visual continuity at the point when a participant walked though the door.

6. Experiment 3

In Experiments 1 and 2, participants' navigation in Building B was affected by the type of connection they had used in Building A. In both of the experiments, six participants used doors in both Building A and the first Building B, and another six participants used hyperlinks in both of these buildings. The largest differences between the two types of connection were likely to occur in the first Building B because participants had become familiarized with navigating VEs (in Building A) but had performed the full test with only one type of connection. Experiment 3 investigated the effects of visual and movement continuity by comparing the navigation of participants who used ghost-doors versions of Building A and Building B, with the 12 participants from Experiment 2 who had used a single type of connection in the first two buildings.

6.1. Method

6.1.1. Participants

Six participants (2 men and 4 women) took part in the experiment. All the participants were either students or graduates, volunteered for the experiment, were paid an honorarium for their participation and had not taken part in the previous experiments. Their ages ranged from 19 to 25 years (M = 22.0).

6.1.2. Materials

Experiment 3 used the same virtual buildings and software as Experiments 1 and 2. The only change for Experiment 3 was the implementation of the ghost-like doors (see above). As in Experiments 1 and 2, participants could not walk through the walls of the rooms.

6.1.3. Procedure

The procedure was the same as for Experiments 1 and 2, but with the exception of the sequence of connection types that participants used in the VEs. All the participants in Experiment 3 repeatedly navigated a ghost-doors version of Building A. Then they repeatedly navigated a ghost- doors version of either Building B1 (three participants) or Building B2 (the other three participants). Finally they navigated a time-delay hyperlinks version of the other Building B. Participants navigated the second Building B so that the procedure and amount of time was similar to Experiment 2.

6.2. Results

To avoid any interaction between the different types of connection, only the data for the first Building B were analyzed. Data for participants in Experiment 3 (ghost-doors) were compared with data for half the participants in Experiment 2 (normal-doors and time-delay hyperlinks). The number of rooms that participants entered was analyzed using repeated measures ANOVAs, with the type of connection treated as a between-groups factor.

There was no effect of connection type on the number of rooms that participants entered in their any-order searches, $F(2, 15) = 1.65$, $p = .23$, but there was an effect of search number, $F(1, 15) = 5.13$, $p = .04$ (see **Table 5**). In the repeated-route searches there were significant effects of connection type, $F(2, 15) = 4.05$, $p = .04$, and search number, $F(5, 15) = 3.93$, $p < .01$ (see **Figure 11**). Planned contrasts showed that the difference between normal- and ghost- doors was significant, $F(1, 15) = 5.47$, $p = .03$, but the difference between ghost-doors and hyperlinks was not significant, $F(1, 15) = 0.06$, $p = .82$. In the short-cut searches there was no effect of connection type, $F(2, 15) = 0.76$, $p = .48$, or search number, $F(1, 15) = 2.82$, $p = .11$, but there was a significant interaction, , $F(2, 15) = 5.33$, $p = .02$.

6.3. Discussion

For all three types of search, the mean number of rooms entered by participants in the ghost-doors group was greater than the mean for the normal-doors group and less than the mean for the hyperlinks group. Participants in the ghost-doors and normal-doors group both experienced the same continuity of movement (they experienced δ-place), but participants in the former group experienced zero visual continuity at the moment in time when they moved from one room to the next, as did participants in the hyperlinks group. The repeated-route data suggest that it is the break in visual continuity, rather than the lack of movement continuity, which has the greater effect on navigation between places that are connected by hyperlinks.

Experiments 1, 2 and 3 investigated the effects of hyperlinks on navigation in VEs that were otherwise conventional (the hyperlinks connected spatially-adjacent rooms). Experiment 4 investigated the effects of "useful" hyperlinks that connected otherwise disparate regions of virtual buildings and allowed participants to travel more efficiently between some rooms. Two different types of hyperlink were used. The first was identical to the type of hyperlink that was used in Experiment 2 (visually discontinuous). The second type was visually continuous, had the appearance of a door, and was chosen because the data from Experiment 3 highlighted the benefits on people's navigation in VEs of preserving visual continuity.

7. Experiment 4

7.1. Method

7.1.1. Participants

Twenty-four participants (6 men and 18 women) took part in the experiment. All the participants were either students or graduates, volunteered for the experiment, were paid an honorarium for their participation and had not taken part in the previous experiments. Their ages ranged from 18 to 34 years (M = 19.1). Each participant was randomly assigned to one of three groups (normal-doors, discontinuous-HLs, or continuous-HLs). One participant asked to withdraw after completing their searches in Building C and was replaced in the experiment.

7.1.2. Materials

Experiment 4 used the same software as the previous experiments. Three versions of two different buildings were created, the layouts of which are shown in **Figures 12** and **13**. In all three versions of each VE the objects were turned on and off when participants entered and left the rooms.

In one version the connections were doors. In the other versions, either two (Building C) or three (Building D) extra connections were added to connect together rooms that were not spatially adjacent. In the second version of each building, all the connections were the same as the hyperlinks that were used in Experiment 2 (visually discontinuous and containing a time-delay). In the third version of each building, all the connections between adjacent rooms were doors, but the extra connections were visually continuous hyperlinks. These hyperlinks had the same appearance as the doors and were used in the same way as the doors.

The effect of traversing one of the visually continuous hyperlinks is best explained using an example. In Building D, a participant who was in the room that contained the fork could walk through a door which lay in the direction of the helicopter but, instead of entering that room, they entered the room that contained the clock, coming from the direction of the cup! This meant that the helicopter and clock appeared to share the same position in Euclidean space, as did the fork and cup. Another consequence of the visually-continuous hyperlinks was that objects like the clock appeared to exist in more than one place (adjacent to the fork and the cup), even though participants were explicitly told there was only one instance of each object. All the rooms were defined using four dimensions (XYZ and a visual zone), and the VE software used the zones to render the correct rooms to the display for each graphics frame. At the moment a participant crossed the threshold of the hyperlink that connected the fork to the clock, the VE software made a sudden change to the participant's actual position in the VE, from the edge of the room that contained the fork to the edge of the room that contained the clock. This change in the participant's actual Euclidean position was the same for the visually continuous and visually discontinuous hyperlinks. However, in the former, only a small change took place in the scene that was shown on the display and that change was no different in magnitude to the changes that took place when the participant moved between spatially-adjacent rooms (hence, visual continuity). When a participant traversed a visually discontinuous hyperlink, the scene changed completely.

7.1.3. Procedure

The procedure was similar to those used in the previous experiments and each participant took approximately 1.5 hr. to complete the experiment. Building C was used to allow participants to practice the interface controls and then perform a practice test. The full test was carried out in Building D. Participants used the same type of connection in both buildings.

For the practice, a participant repeatedly navigated Building C. Once they had practiced the controls they performed an any-order search for the two target objects (see Figure 12). Then they re-entered the VE and performed five repeated-route searches for the pair. Each search started at the flower and ended at the knife and, as in the previous experiments, the display was blanked for 5 s when the knife had been selected. Finally, the participant was asked to draw a sketch map.

The full test was performed in Building D. The participant performed two any-order searches and then several pairs searches. In the pair searches, half of the participants in each group started at the house, searched for the spanner, waited while the screen was blank, started at the saucepan, searched for the helicopter, and so on until they had searched for each pair a total of five times. The other participants searched for the helicopter and then the spanner. Finally, the participant was asked to draw a sketch map.

As in the VEs used in the previous experiments, barriers to movement (walls without connections) were included so that each target pair was connected by a specific shortest route. The shortest route from the saucepan to the helicopter involved entering three rooms (the starting room was not counted) which were connected by links that were present in all three versions of the VE. The spanner could be reached from the house by entering three rooms in the continuous-HLs and discontinuous-HLs versions of the VE but, in the normal-doors version, a participant had to enter at a minimum of eight rooms.

7.2. Results

As in the other experiments, only the data for the full test (Building D) will be reported. This was done by comparing the number of rooms that participants entered using the same types of repeated measures ANOVAs as were used in the previous experiments. All of the participants in the continuous-HLs and discontinuous-HLs groups used the extra connections.

In the any-order searches, participants in the continuous-HLs group ($\underline{M}$ = 32.7) entered fewer rooms than participants in the normal-doors group ($\underline{M}$ = 50.3) or the discontinuous-HLs group ($\underline{M}$ = 40.6). The main effect for these searches was only marginally significant, $\underline{F}(2, 21) = 3.08$, $\underline{p} = .07$, but planned contrasts showed that the difference between the continuous-HLs and the normal-doors groups was significant, $\underline{F}(1, 14) = 6.16$, $\underline{p} = .02$. The other inter-group differences were not significant.

The two target object pairs were analyzed separately because the shortest route between the house and spanner was larger in the doors condition (8 rooms) than in the hyperlink conditions (3 rooms). As expected, repeated measures ANOVA for the house/spanner pair showed that there was a main effect of group, $F(2, 21) = 13.72$, $p < .01$, but there was no effect of search number, $F(4, 21) = 1.68$, $p = .16$ (see **Figure 14**). Planned contrasts showed that there were significant differences between the continuous-HLs and normal-doors groups, $F(1, 14) = 21.85$, $p < .01$, and the discontinuous-HLs and normal-doors groups, $F(1, 14) = 19.21$, $p < .01$, but not between the two hyperlinks groups, $F(1, 14) = 0.08$, $p = .77$. A repeated measures ANOVA for the saucepan/helicopter pair showed that there was no main effect of group, $F(2, 21) = 0.33$, $p = .73$, or of search number, $F(4, 21) = 1.02$, $p = .40$. Averaged across searches, group means for the normal-doors, continuous-HLs and discontinuous-HLs were 11.5, 11.0, and 13.7, respectively.

7.2.1. Sketch Maps

The cells and spatial names metrics would have been difficult to calculate because participants in the continuous-HLs and discontinuous-HLs conditions could, legitimately, have drawn some of the rooms in many different positions. For this reason, the sketch maps were analyzed using only the object names, adjacent rooms and links metrics. There were no effects of connection type on any of the metrics and the mean data, averaged across all three groups, were 44% (object names), 12% (adjacent rooms) and 16% (links).

7.2.2. Random Walk Analysis

The connectivity of the rooms in Building D was different for participants in the normal-doors group than for participants in the hyperlinks groups. For this reason, we calculated the number of rooms that would be entered in the any-order searches by 1,000,000 random walks. In the normal-doors building, participants had to enter a minimum of 16 rooms and 50% of their random walks were completed by a hypothetical participant who visited less than 46 rooms. In the hyperlinks buildings the minimum and 50% figures were 13 and 47 rooms, respectively. The gap between the normal-doors and hyperlinks then widened slightly as the percentage of random walks increased.

7.3. Discussion

Both the present, and earlier, studies have shown that people tend to be very disoriented when they initially navigate a VE (i.e., they make more random choices of route). The random walk data for Experiment 4 indicated that participants in the hyperlinks groups were likely to enter more rooms than the normal-doors participants in the early stages of the experiment (the any-order searches). In fact, the opposite occurred and this was both interesting and unexpected. The most likely explanation is that the extra connections allowed participants to travel directly from one part of the VE to another, immediately changing the region in which they searched. However, it is not known whether participants chose to do this deliberately, or by chance. Whatever the explanation, the data show that participants found the extra connections useful for the any-order searches.

The data from the two repeated-route searches and the sketch maps indicate that participants had great difficulty learning routes through, and the overall layout of, the VE. The probable reason for this is the increase in the number of rooms and connections in Building D when compared with the VEs used in the previous experiments, although 25 rooms is still a very modest size.

In the house/spanner searches, participants in both hyperlinks groups exploited the extra connections, which allowed them to jump from one region of the VE to another, again confirming the usefulness of such a feature. However, none of the groups improved significantly with successive searches and, instead, seemed unable to learn the shortest route. This was not particularly surprising for the normal-doors group, for whom even the shortest route involved traveling through eight rooms, but the other groups only had to learn a route through three rooms, the same as had successfully been learned in the previous experiments. In the saucepan/helicopter searches not even the normal-doors group, who navigated a conventionally-connected VE, improved significantly with successive searches.

8. General Discussion

The experiments show that hyperlinks affect the ability of participants to navigate in VEs, but the effect is modest when compared to the general difficulty participants had navigating conventional VEs in both this and previous studies. In the present study, when participants could traverse hyperlinks instantaneously they changed their behavior and searched for the objects in a more random manner, with a combined effect of entering more rooms, but finding the objects more quickly, than in a conventionally-connected VE. However, the VEs were fairly small (a maximum of 25 rooms) and a random search becomes less effective as the complexity of a VE increases. When the effect of traversal time was eliminated by providing a time-delay, the hyperlinks were still found to inhibit participants' navigation, especially when they first entered a VE.

Two other differences between conventional and hyperlink movement though VEs occur in the amounts of visual and movement continuity that are present. The introduction of momentary visual discontinuity, through the use of ghost-doors, suggested that it was the visual component of continuity which was the most important.

In Experiment 4 participants exploited the extra connections, which allowed them to jump from one part of the VE to another, to increase the efficiency of their navigation. This opens the way for hyperlinks to be used in a wide variety of VE applications, in the same way as they are already well-established in written documentation that is stored as hypertext. If the absolute minimum of disruption to navigation is desired then hyperlinks should be designed so that they preserve visual continuity. Unfortunately, the implementation of visually continuous hyperlinks requires changes to VE rendering software because the z-buffer algorithm used in most 3-D computer graphics applications (see Foley, van Dam, Feiner, & Hughes, 1990) only works with non-overlapping, Euclidean spaces. An alternative solution, which would provide partial visual continuity, would be to provide a window at each end of a hyperlink that contained a picture of the scene at the other end.

Questions remain about the type of mental representations (e.g., cognitive map; Tolman, 1948) that people form when they navigate. It may, or may not, be a map (see Kitchin, 1994) but it does have some map-like properties. For example, Thorndyke and Hayes-Roth (1982) showed that a person who learned the layout of a building by navigation eventually developed survey-type knowledge (straight-line distances and relative directions) that was at least as accurate as that of a person who learned the knowledge from a map which presented the information directly, and Ruddle et al. (1997) found a similar trend occurred in a VE.

A person's mental representation of a real-world space sometimes contains systematic distortions that would produce discontinuities if directly mapped to Euclidean space. For example, in the study by Tversky (1981), participants tended to draw sketch maps in which streets intersected at 90°, even when that was not the case, and non-parallel streets were made parallel. Sadalla and Magel (1980) and Sadalla and Staplin (1980) showed that changes of direction and the number of path intersections influence the distance that a person perceives they travel and, when a person estimates the straight-line distance between two places, their estimates in one direction are sometimes different to their estimates in the other direction, breaking one of Euclid's axioms (see Baird, Wagner, & Noma, 1982).

Environments that contain hyperlinks are likely to produce further modifications to a person's spatial representations. If hyperlinks are implemented using hybrid integration then a person's representation may be the same as their representation of a conventional space but with the hyperlinks stored as either extra routes, or routes that allow particularly rapid movement. With island integration, where hyperlinks provide the only mechanism for movement between different regions, one possibility is that a person's representation may consist solely of route information (effectively a network of nodes and links), but people might be influenced toward forming a conventional representation if they know a VE is a model of a real place (e.g., Virtual Tourism: Paris, 1996).

The primary challenge for designers is to make VEs that contain hyperlinks more navigable, while maintaining the advantages of discontinuous movement. Techniques that aid navigation in conventional VEs may also dramatically reduce learning time when hyperlinks are used. Techniques that are being investigated for use with maps of sections of the WWW (e.g. Mukherjea & Foley, 1995; Mukherjea, Foley & Hudson, 1995; Zizi & Beaudouin-Lafon, 1995) may be combined with map displays developed from the findings of Darken and Sibert (1993, 1996) and Ruddle et al. (1999a) to provide effective navigational aids for all types of VE.

Finally, the present study provides a performance baseline for studies of navigation in VEs that contain hyperlinks. Further experimental studies are required to investigate the other attributes that were identified in this article.

Acknowledgements

This work was supported by grant GR/L 28449 from the Engineering and Physical Sciences Research Council. We are grateful to Dr. Patrick Péruch for his helpful comments on an earlier draft.

References

Arthur, K. W., Booth, K. S., & Ware, C. (1993). Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems*, **11**, 239–265.

Baird, J. C., Wagner, M., & Noma, E. (1982). Impossible cognitive spaces. *Geographical Analysis,* **14**, 204-216.

Bliss, J. P., Tidwell, P. D., & Guest, M. A. (1997). The effectiveness of virtual reality for administering spatial navigation training to firefighters. *Presence: Teleoperators and Virtual Environments,* **6**, 73-86.

Bowman, D. A., Koller, D., & Hodges, L. F. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Proceedings of the Virtual Reality Annual International Symposium (VRAIS'97)*, pp. 45-52.

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.

Conklin, J. (1987). Hypertext: An introduction and survey. *Computer,* **20**, 17-41.

Darken, R. P., & Sibert, J. L. (1993). A toolset for navigation in virtual environments. *Proceedings of ACM User Interface Software & Technology (UIST '93)*, pp. 157-165.

Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction,* **8**, 49-71.

Evans, G. W., Skorpanich, M. A., Gärling, T., Bryant, K. J., & Bresolin, B. (1984). The effects of pathway configuration, landmarks and stress on environmental cognition. *Journal of Environmental Psychology,* **4**, 323-335.

Foley, J. D., van Dam, A., Feiner, S. K., & Hughes, J. F. (1990). *Computer Graphics: Principles and Practice* (2nd ed.). Reading, MA: Addison-Wesley.

Furnas, G. W. (1997). Effective View Navigation. *Proceedings of Computer Human Interfaces Conference (CHI'97)*, pp. 367-374.

Grant, S. C., Magee, L. E. (1998). Contributions of proprioception to navigation in virtual environments. *Human Factors,* **40**, 489-497.

Hancock, P. A., Hendrix, C., & Arthur, E. (1997). Spatial mental representations in virtual environments. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*, pp. 1143-1147.

Kim, H., & Hirtle, S. C. (1995). Spatial metaphors and disorientation in hypertext browsing. *Behaviour and Information Technology,* **14**, 239-250.

Kitchin, R. M. (1994). Cognitive maps: What are they and why study them?. *Journal of Environmental Psychology,* **14**, 1-19.

Lynch, K. (1960). *The image of the city*. Cambridge: MIT Press.

May, M., Péruch, P., & Savoyant, A. (1995). Navigating in a virtual environment with map-acquired knowledge: Encoding and alignment effects. *Ecological Psychology*, **7**, 21-36.

Mukherjea, S., & Foley, J. D. (1995). Showing the context of Nodes in the World-Wide Web. *Proceedings of Computer Human Interfaces Conference (CHI'95)*, pp. 326-327.

Mukherjea, S., Foley, J. D., & Hudson, S. (1995). Visualizing complex hypermedia networks through multiple hierarchical views. *Proceedings of Computer Human Interfaces Conference (CHI'95)*, pp. 331-337.

O'Hara, K. P., & Payne, S. J. (1998). The effects of operator implementation cost on planfulness of problem solving and learning. *Cognitive Psychology,* **35**, 34-70.

Richard, P., Birebent, G., Coiffet, P., & Langrana, N. (1996). Effect of frame rate and force feedback on virtual object manipulation. *Presence: Teleoperators and Virtual Environments*, **25**, 95–108.

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.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1998). Navigating large-scale “desk-top” virtual buildings: Effects of orientation aids and familiarity. *Presence: Teleoperators and Virtual Environments,* **7**, 179-192.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999a). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied,* **5**, 54-75.

Ruddle, R. A., Payne, S. J., & Jones, D. M. (1999b). Navigating Large-Scale Virtual Environments: What Differences Occur Between Helmet-Mounted and Desk-Top Displays? *Presence: Teleoperators and Virtual Environments,* **8**, 157-168.

Sadalla, E. K., & Magel, S. G. (1980). The perception of traversed distance. *Environment and Behavior*, **12**, 65-79.

Sadalla, E. K., & Staplin, L. J. (1980). The perception of traversed distance: Intersections. *Environment and Behavior*, **12**, 167-182.

Stanton, D., Wilson, P., & Foreman, N. (1996). Using virtual reality environments to aid spatial awareness in disabled children. *Proceedings of the 1st European Conference on Disability, Virtual Reality and Associated Technology*, pp. 93-101.

Thorndyke, P. W., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology,* **14**, 560-589.

Tlauka, M., & Wilson, P. N. (1994). The effects of landmarks on route-learning in a computer-simulated environment. *Journal of Environmental Psychology,* **14**, 305-313.

Tlauka, M., & Wilson, P. N. (1996). Orientation–free representations from navigation through a computer–simulated environment. *Environment and Behavior*, **28**, 647-664.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, **55**, 189-208.

Tversky, B. (1981). Distortions in memory for maps. *Cognitive Psychology,* **13**, 407-433.

Utting, K., & Yankelovich, N. (1989). Context and orientation in hypermedia networks. *ACM Transactions on Information Systems,* **7**, 58-84.

Virtual Tourism: Paris. (1996). Cupertino, CA: Apple Computer. [Computer Software].

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. New Jersey: Erlbaum.

Wickens, C. D. (1992). *Engineering Psychology and Human Performance* (2nd ed.). New York: HarperCollins.

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.

Woods, D. D. (1984). Visual momentum: A concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies*, **21**, 229–244.

Zizi, M., & Beaudouin-Lafon, M. (1995). Hypermedia exploration with interactive dynamic maps. *International Journal of Human-Computer Studies,* **43**, 441-464.

Table 1.

Attributes of hyperlinks that may affect navigation in VEs.

Attribute	Ease or speed of navigation		
	Less		More
Integration	island	--	hybrid
Appearance	object	link menu	hot-list
Time	fixed	distance-related	instantaneous
Motion profile	jump	--	variable velocity
Preview	none	single-step	signpost
Visual continuity	0%	--	100%
Directionality	unidirectional	--	bidirectional
Directional affordance	inconsistent	indeterminate	consistent
Rotational offset	180°	--	0°
Crossings	non-planar	--	planar

Table 2

Number of rooms entered in the any-order and short-cut searches in Experiment 1

Connection	Any-order				Short-cut			
	Search 1		Search 2		Search 1		Search 2	
	M	SE	M	SE	M	SE	M	SE
Doors	28.3	2.4	24.8	2.5	8.9	0.9	7.6	0.7
Hyperlinks	33.8	3.2	30.2	3.0	12.4	1.1	10.2	0.9

Table 3

Time taken in the any-order and short-cut searches in Experiment 1

Connection	Any-order				Short-cut			
	Search 1		Search 2		Search 1		Search 2	
	M	SE	M	SE	M	SE	M	SE
Doors	206	15.7	128	16.6	82	8.9	63	6.2
Hyperlinks	349	35.6	214	26.6	51	3.8	41	3.4

Table 4

Number of rooms entered in the any-order and short-cut searches in Experiment 2

Connection	Any-order				Short-cut			
	Search 1		Search 2		Search 1		Search 2	
	M	SE	M	SE	M	SE	M	SE
Doors	28.0	3.0	24.5	1.9	8.5	0.9	7.3	0.7
Hyperlinks	35.0	4.2	24.8	2.2	8.9	1.1	8.3	0.8

Ruddle, R. A., Howes, A., Payne, S. J., & Jones, D. M. (2000) The effects of hyperlinks on navigation in virtual environments. *International Journal of Human Computer Studies, 53*, 551-581.

Table 5

Number of rooms entered in the any-order and short-cut searches in first Building B in Experiment 3

Connection	Any-order				Short-cut			
	Search 1		Search 2		Search 1		Search 2	
	M	SE	M	SE	M	SE	M	SE
Doors	28.6	7.9	23.3	2.7	8.3	1.9	5.5	1.0
Ghost	40.1	7.7	28.3	2.1	8.7	1.7	10.9	2.4
Hyperlinks	44.4	11.6	29.9	5.3	9.0	1.9	7.5	2.4

Figure Captions

Figure 1. Virtual buildings that illustrate hybrid and island integration. Rooms are shown shaded and corridors are shown white. The only method of travelling between the two islands is by traversing hyperlink H_3 or H_4.

Figure 2. A virtual building that contains two hyperlinks, H_b and H_c, that respectively have consistent and inconsistent directional affordance

Figure 3. A virtual building that contains hyperlinks with 0° and 90° rotational offsets. The dashed lines indicate a person's direction of view at each end of the hyperlinks.

Figure 4. Plan view of Building A. The any-order searches started in the room that contains the underlined object name. The object start/finish pairs are shown in parentheses. The gaps in the walls indicate the position and connectivity of the doors and hyperlinks.

Figure 5. Plan view of Buildings B1 (left) and B2 (right). The any-order searches started in the room that contains the underlined object name. The object start/finish pairs are shown in parentheses. The gaps in the walls indicate the position and connectivity of the doors and hyperlinks.

Figure 6. An interior view of Building B1. The hyperlinks (the Hs) lead to the truck (left) and traffic light (right).

Figure 7. Number of rooms that participants entered in their repeated-route searches in Experiment 1. Error bars indicate standard error of the mean.

Figure 8. The time that participants took to complete their repeated-route searches in Experiment 1. Error bars indicate standard error of the mean.

Figure 9. Number of rooms predicted by the random walk program, and entered by participants, in the first any-order searches in Experiment 1.

Figure 10. Number of rooms that participants entered in their repeated-route searches in Experiment 2. Error bars indicate standard error of the mean.

Figure 11. Number of rooms that participants entered in their repeated-route searches in the first Building B for Experiment 3. Error bars indicate standard error of the mean.

Figure 12. Plan view of Building C. The any-order searches started in the room that contains the underlined object name. The object start/finish pairs are shown in parentheses. The gaps in the walls indicate the position and connectivity of the doors and hyperlinks. The arrows indicate the rooms that were connected by extra connections.

Figure 13. Plan view of Building D. The any-order searches started in the room that contains the underlined object name. The object start/finish pairs are shown in parentheses. The gaps in the walls indicate the position and connectivity of the doors and hyperlinks. The arrows indicate the rooms that were connected by extra connections.

Figure 14. Number of rooms that participants entered in their repeated-route searches for the house-spanner pair in Experiment 4. Error bars indicate standard error of the mean.

Figure 1

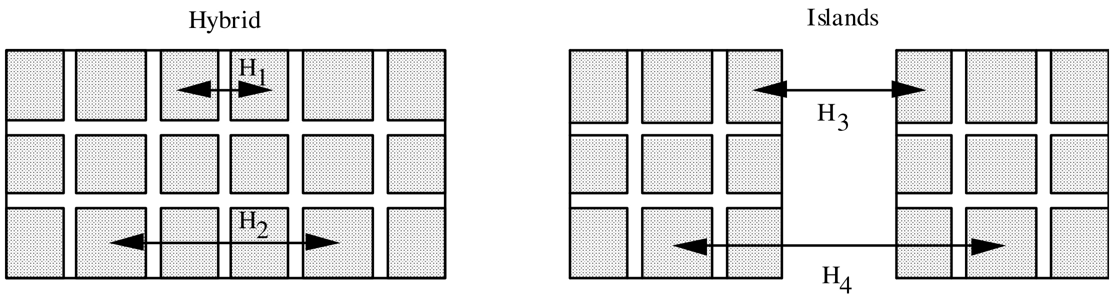

Ruddle, R. A., Howes, A., Payne, S. J., & Jones, D. M. (2000) The effects of hyperlinks on navigation in virtual environments. *International Journal of Human Computer Studies, 53*, 551-581.

Figure 2

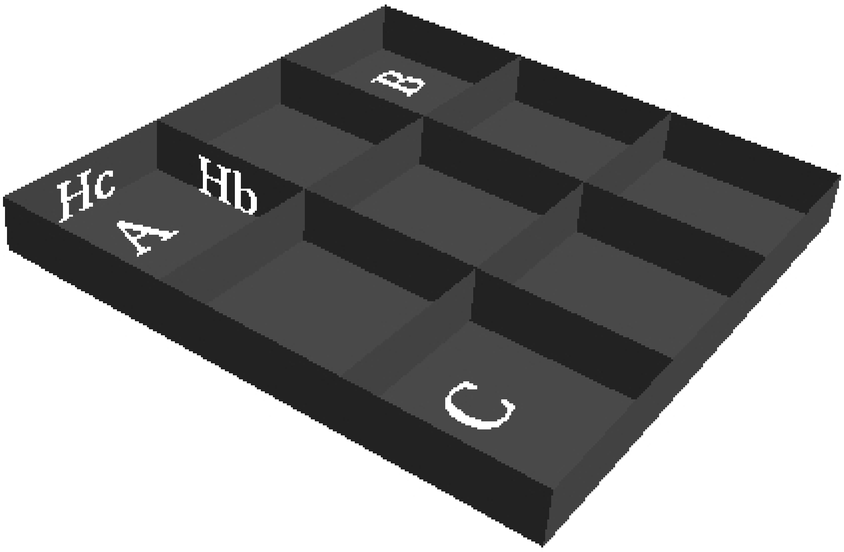

Figure 3

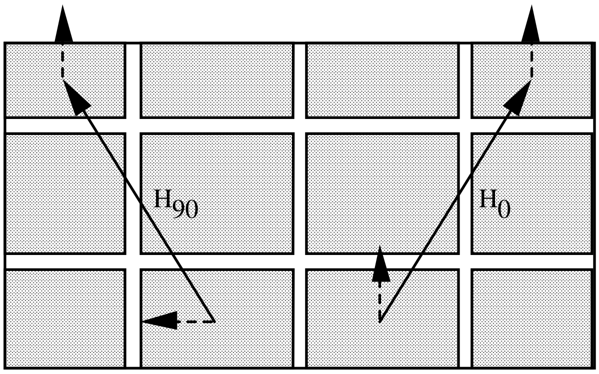

Figure 6

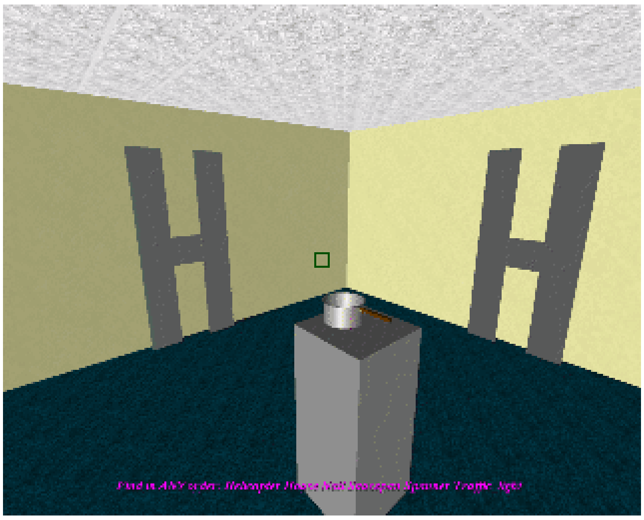

Ruddle, R. A., Howes, A., Payne, S. J., & Jones, D. M. (2000) The effects of hyperlinks on navigation in virtual environments. *International Journal of Human Computer Studies, 53*, 551-581.