Available online at www.sciencedirect.com
ScienceDirect
Int. J. Human-Computer Studies 67 (2009) 237–255
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

Multimodal selection techniques for dense and occluded 3D virtual environments ☆

Lode Vanacken[a,*], Tovi Grossman[b], Karin Coninx[a]

[a]Hasselt University – tUL – IBBT, Expertise Centre for Digital Media, Wetenschapspark 2, B-3590 Diepenbeek, Belgium
[b]Autodesk Research, Autodesk, 210 King Street East, Toronto, Canada

Received 21 April 2008; received in revised form 22 August 2008; accepted 2 September 2008
Available online 20 September 2008

Abstract

Object selection is a primary interaction technique which must be supported by any interactive three-dimensional virtual reality application. Although numerous techniques exist, few have been designed to support the selection of objects in dense target environments, or the selection of objects which are occluded from the user's viewpoint. There is, thus, a limited understanding on how these important factors will affect selection performance. In this paper, we present a set of design guidelines and strategies to aid the development of selection techniques which can compensate for environment density and target visibility. Based on these guidelines, we present new forms of the ray casting and bubble cursor selection techniques, which are augmented with visual, audio, and haptic feedback, to support selection within dense and occluded 3D target environments. We perform a series of experiments to evaluate these new techniques, varying both the environment density and target visibility. The results provide an initial understanding of how these factors affect selection performance. Furthermore, the results showed that our new techniques adequately allowed users to select targets which were not visible from their initial viewpoint. The audio and haptic feedback did not provide significant improvements, and our analysis indicated that our introduced visual feedback played the most critical role in aiding the selection task.

Keywords: Virtual reality; 3D interaction; Selection; Multimodal feedback

1. Introduction

In recent years, three-dimensional (3D) display technologies, such as immersive virtual reality (VR) systems (Buxton and Fitzmaurice, 1998), or non-immersive fish-tank VR systems using LCD shutter stereo-glasses (Ware et al., 1993), have significantly improved in display quality. Experimental evaluations have also shown that these displays can improve the user's ability to perceive virtual 3D scenes (Ware and Franck, 1996), making them a potentially beneficial alternative for 3D applications.

One of the primary techniques which must be supported in any interactive application which must be supported is object selection. Within the realm of 3D virtual environments, selection has been repeatedly identified as one of the fundamental tasks (Mine, 1995; Bowman et al., 2004). However, when selecting objects in a 3D environment, the standard 2D mouse metaphor breaks down, as the targets will have 3D coordinates, which the user must somehow specify. As such, it is important for VR researchers to consider new selection techniques, specifically designed for 3D environments.

Indeed, research in VR environments has introduced numerous techniques for object selection. Most commonly seen are hand extension techniques (Mine, 1995), for which the 3D coordinates of the hand or handheld input device are mapped to the 3D coordinates of a virtual cursor, and ray casting techniques (Liang and Green, 1994), for which a virtual ray is cast into the scene, and made to intersect targets of interest. Despite the numerous designs

☆Extension of "Exploring the effects of environment density and target visibility on object selection in 3D virtual environments", by L. Vanacken, T. Grossman, and K. Coninx which appeared in the proceedings of the IEEE Symposium on 3D User Interfaces, pp. 115–122 © 2007 IEEE.

*Corresponding author. Tel.: +32 011268411; fax: +32 011268499.
E-mail addresses: lode.vanacken@uhasselt.be (L. Vanacken), tovi.grossman@autodesk.com (T. Grossman), karin.coninx@uhasselt.be (K. Coninx).

1071-5819/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ijhcs.2008.09.001

and evaluations of such techniques, there are a number of important factors which remain to be fully understood, two of which we focus on in this paper.

The first is the density of targets in the environment. This can greatly affect the performance of a selection technique, particularly in a 3D environment. For example, with traditional ray casting techniques, only the first intersected target can be selected, making it difficult to select targets behind dense areas (Liang and Green, 1994). Alternatively, allowing multiple targets to be intersected, either along the length of the ray (Hinckley et al., 1994), or in some cases, within its conic selection area (Liang and Green, 1994), results in an ambiguity of the user's intended target. Recently, researchers have attempted to address this problem of ambiguity with new selection techniques (Olwal and Feiner, 2003; Grossman and Balakrishnan, 2006; Wyss et al., 2006), but the effect of target density on such techniques remains to be explored.

The second factor which we investigate in this paper is the visibility of the goal target. While, in many cases, users' intended targets are visible, there may be cases where the intended targets are occluded from their view by another object in the scene. Generally, users are required to either rotate the scene or to switch to a different viewing mode, to allow them to select their desired target. These extra steps may be time consuming, requiring the addition of modes and buttons to the interface. It is, therefore, of interest to develop techniques which allow users to seamlessly select targets which are occluded from their view, and to understand how the visibility of the target affects selection performance.

In the following, we investigate various candidate selection techniques, which may allow for efficient selection of targets in both sparse and dense 3D environments. Furthermore, we augment these techniques with various forms of multimodal feedback to also allow for the selection of targets which are not visible from the user's viewpoint. In an initial experiment, we augment a depth ray and 3D bubble cursor with only visual feedback, and compare these techniques to a standard point cursor, also augmented with visual feedback to allow for the selection of occluded targets. The results show that both the depth ray and the 3D bubble cursor outperform the point cursor in all conditions, with the depth ray performing best overall. Furthermore, the additional visual feedback which we introduced to all techniques allowed users to adequately select targets which were not visible from their initial viewpoint. However, our observations indicated that the visual feedback was not sufficient to completely overcome the difficulties associated with selecting targets in dense and occluded environments. This visual feedback stimulates only the (already heavily loaded) human visual system, leaving the powerful human senses of touch and hearing unused. Such other forms of feedback have been suggested as a way to improve interaction and reduce the load on any one sense (Bolt, 1980; Oviatt, 2002). Therefore, in a second experiment, we augment the selection techniques with audio and haptic feedback, in addition to the visual feedback used in the first experiment. Results show that while participants did prefer the presence of the multimodal feedback, the quantitative advantages were only slight, and non-significant. These results indicate that providing adequate visual feedback is most critical for selections in dense and occluded 3D environments.

This paper is an extension of our earlier research presented at *IEEE 3D User Interfaces 2007* (Vanacken et al., 2007). It extends the earlier work with the development of novel haptic, audio, and visual feedback augmentations to the previously evaluated selection techniques. These augmentations are then evaluated in an additional study, and the results are compared to the initial study.

2. Related work

Researchers in 3D virtual environments have often categorized selection as one of the four basic interactions (along with navigation, manipulation, and data input) (Mine, 1995; Bowman et al., 2004). Selection allows users to specify an object, with which they wish to manipulate or interact. Because it is such a critical task, there has been a wide variety of research looking into various techniques for supporting it. In this section, we first provide a review of selection techniques for 3D environments. We then discuss selection techniques which have been studied under dense target environments, followed by a discussion of techniques for overcoming target occlusions. Lastly, we discuss the use of multimodal feedback, and its application to selection techniques.

2.1. 3D selection techniques

One of the earliest implementations of selection for 3D environments was Liang and Green's (1994) "laser gun" ray casting technique. With this technique, a ray is emitted from the user's hand, so the user has control over the origin and trajectory of the ray, much like using a physical laser pointer. One observed problem with this technique was that it was difficult to select distant and small objects due to the angular accuracy which was required. To overcome this, they introduced a technique called "spotlight selection" where, instead of emitting a ray, the user emits a conic selection area, originating from the user's hand. Since this original work, there have been a number of iterations on the ray casting metaphor in the 3D research community, such as aperture based selection (Forsberg et al., 1996) and 2D image plane selection (Pierce et al., 1997). However, a general problem with the ray casting metaphor is that it can be difficult to select a target that is behind a dense surrounding of targets, since with traditional ray casting only the first target to be intersected is selected. Diverging from this traditional implementation, and allowing multiple targets to be intersected by the ray, or increasing the activation area to a conic selection area, creates an ambiguity of the intended target. Existing strategies to address this ambiguity are discussed in the next section.

The most common alternative to the ray casting metaphor is the hand extension metaphor, where the user controls the X, Y, and Z coordinates of a 3D cursor (Hinckley et al., 1994; Mine, 1995; Poupyrev et al., 1996). Mine (1995) states that in local interactions, a direct mapping from the user's hand to a 3D "virtual cursor or drone" could be used to select an object. For distant objects, the go–go technique explores the use of nonlinear mappings between the user's hand and 3D cursor, extending the range which the cursor can cover (Poupyrev et al., 1996). One drawback of the hand extension metaphor is that the selections are constrained by three dimensions, resulting in longer selection times (Poupyrev et al., 1998; Bowman et al., 1999; Grossman and Balakrishnan, 2004, 2006; De Boeck et al., 2006). As an alternative to a simple 3D point cursor, Zhai et al. (1994) developed the silk cursor, which is a semi-transparent 3D volume cursor. Using a volume increases the activation area of the cursor, and it was shown to increase performance in a 3D target tracking task. While a volume cursor could reduce target acquisition times, it, once again, produces difficulty when interacting in dense target environments, as multiple targets may fall within the boundaries of the cursor's volume.

2.2. Selection techniques for dense environments

For ray casting techniques, generally only the first intersected object will be selected, even though the ray can intersect multiple objects simultaneously. Under this implementation, it may be difficult or even impossible to select objects that are further away, depending on the density of targets. While research has shown that users may choose to navigate to a new position such that the desired object is closer or not occluded (Ware and Lowther, 1997), there may be cases where the target will be occluded regardless of the viewpoint. Furthermore, it may be desirable for the user to have the ability to perform the selection in-place, without having to switch to and perform navigation.

To address this, Liang and Green (1994) developed a metric for the spotlight selection to determine which object would be selected when multiple targets were captured, based on the distance between the target to the apex and central axis of the cone. An interesting extension to spotlight selection is shadow cone selection (Steed and Parker, 2004), which selects targets by sweeping out an area with a cone selection cursor. Other metrics have also been proposed (Steed, 2006), but it is unclear how well they will work in dense target environments, as they have not been formally evaluated.

An alternative to defining these predictive metrics is to provide an explicit mechanism for users to specify their intended target among all those intersected (Hinckley et al., 1994). Grossman et al. (2004) used forwards and backwards hand movements to cycle through intersected objects, however, limited visual feedback was provided to the user. Olwal and Feiner (2003) describe the flexible pointer, which allows users to bend the cursor to avoid intersecting other targets, but it requires two 6 DOF (degree-of-freedom) devices to control the cursor. Another technique requiring two input devices is iSith (Wyss et al., 2006), where two rays are intersected to define a target location.

In a recent study, Grossman and Balakrishnan (2006) designed and evaluated several new ray casting techniques which allowed for multiple target disambiguation. Of their tested techniques, they found the depth ray to be most successful. The depth ray augments the ray cursor with a depth marker, which can be moved forwards and backwards along the length of the ray, with similar movements of the hand. The intersected target closest to this depth marker is the one which can be selected. Although the study used a dense environment, the environment was constant throughout the experiment, so they were not able to study the effect of the environment density on the technique.

In a study that did vary target densities, Looser et al. (2007) compared three different selection techniques, including ray casting, for tabletop augmented reality. They found no significant difference between the two densities, although, they argued that this might have been because the difference between the densities which they tested was small.

There has also been work in selecting targets in dense two-dimensional (2D) environments. Most notably, the bubble cursor (Grossman and Balakrishnan, 2005) is an area cursor that dynamically changes its size to always capture only the closest target. The technique was studied in environments of varied density, and found to provide efficient selections in sparse and dense environments. Such a technique could be an interesting alternative to a static 3D volume cursor, and we will explore such a technique in this paper.

2.3. Overcoming target occlusion

Some of the above techniques for dense environments explore the scenario of multiple objects occluding the ray cursor's approach to a goal target. However, none explore the issue of a target being completely occluded from the user's viewpoint. Our literature review found limited research on this topic. In 3D desktop applications, users generally rotate the scene so that their target of interest becomes visible. This has also been shown to be an effective approach in 3D virtual environments (Flasar and Sochor, 2007). Another approach is to switch to a viewing mode, such as wireframe rendering, so that all targets become visible. Some more specialized viewing modes have also been studied in augmented reality environments (Livingston et al., 2003). Other techniques to reduce occlusions, such as interactively distorting the space (Carpendale et al., 1997;

McGuffin et al., 2003; Elmqvist, 2005), or morphing the viewing projection (Elmqvist and Tsigas, 2006), have also been explored. Elmqvist and Tsigas (2008) provide a thorough overview of other possible techniques. However, such techniques are generally independent from the selection mechanisms.

Some relevant techniques in the 2D realm are the tumble and splatter techniques (Ramos et al., 2006). These techniques allow users to first spread out layers of objects which are occluding each other in a 2D environment so that they are all visible. The user can then select and perform interactions with any of these objects.

2.4. Adding multimodal feedback

Multimodal output during object selection has been thoroughly studied in 2D (Akamatsu et al., 1995; Cockburn and Brewster, 2005). For 3D virtual environments, possible benefits for multimodal feedback have been found (Wang and MacKenzie, 2000; Unger et al., 2002; Sallnäs and Zhai, 2003) and its use within traditional selection techniques has been explored (Arsenault and Ware, 2000; Vanacken et al., 2006). However, its use within selection techniques that address target densities and occlusions has not been investigated. Multimodal feedback of collision and contact has been used in many studies (Arsenault and Ware, 2000; Unger et al. 2002; Lindeman, 2003), but it is not clear how such techniques can be directly applied to object selection. We now outline the existing literature related to the addition of force and audio feedback to selection.

2.4.1. Force feedback

A large proportion of haptic feedback research concentrates on selecting items, such as icons and menu items, in Graphical User Interfaces (GUIs) (Oakley et al., 2002; Lécuyer et al., 2004; Ahlström, 2005; Ahlström et al., 2006; Smyth and Kirkpatrick, 2006). Oakley et al. (2000) investigated the effects of different kinds of force feedback, such as textures, friction, and "gravity wells". They found gravity wells, a "snap-to" effect which pulls the user to the center of the target, seemed best in improving performance time and error reduction. However, follow up research, which introduced distractor targets to the task, found that when the cursor needed to pass over the distractor targets, the gravity wells became problematic (Hwang et al., 2003; Keuning, 2003; Ahlström, 2005; Ahlström et al., 2006).

In 3D virtual environments, research into the use of haptics for selection is limited. Arsenault and Ware (2000) found that the addition of force feedback in a tapping task improved the user's performance, probably due to the fact that the cursor bounced off a target actually speeding up its progress back to the other target. Wall et al. (2002) investigated 3D gravity wells for ray casting selection. The haptic feedback did improve the accuracy, but not the performance time. Vanacken et al. (2006) augmented a virtual hand and a 2D bubble cursor, used as an aperture selection technique (Forsberg et al., 1996), with haptic feedback. Results found that the haptic feedback resulted in significant speed gains for the virtual hand technique but not for the 2D bubble cursor, possibly due to its dominating visual feedback. Finally, Kim and Kwon (2007) showed that force feedback can help the user become more aware of depth in a 3D environment, using grid planes placed along the Z-plane.

2.4.2. Audio feedback

The addition of sound and its advantages has also been explored in 2D GUIs (Brewster, 1998a). Many different widgets (menus, buttons, scroll bars, progress bars, etc.) have been augmented with non-speech sound, called earcons, resulting in lower error rates and decreased task completion times (Brewster, 1998a, b). Audio feedback has also been used for selection, to indicate that the cursor has reached a target (Akamatsu et al., 1995; Cockburn and Brewster, 2005; Vanacken et al., 2006). Akamatsu et al. (1995) found that such audio feedback did not improve overall selection time, but it did reduce the time spent over the target, as the sound made users react faster. Cockburn and Brewster (2005) found that the addition of audio feedback reduced mean selection times by 4.2%, but that combining sound with other feedback modalities, which improve selection time independently, does not assure the further improvements. In 3D environments, Vanacken et al. (2006) also found that earcons improved reaction time, but not for the 2D bubble cursor, which seemed to give sufficient visual feedback.

2.5. Summary of related work

To summarize our literature review, there is large body of research on selection techniques for 3D environments. However, less research has focused on supporting selection in dense target environments, and even less on the selection of objects which are fully occluded. While some promising techniques do exist, to date there has not been an exploration of how these techniques are affected by the environment density and target visibility. With regard to other modalities of feedback, audio earcons have been shown to improve reaction times when visual feedback is inadequate. Force feedback has also been shown to be a viable technique, but evidence suggests that it may encounter difficulties in dense target environments.

3. Design guidelines and strategies

Before introducing the techniques which we present and evaluate in this paper, we first discuss some high-level design guidelines which we had for these techniques. In this section, we present these design guidelines, followed by our proposed design strategies for satisfying these guidelines.

3.1. Design guidelines

Our literature review of the previous work on selection techniques in 3D environments revealed that the environment density and visibility of the goal target are factors which are not well understood. As such, it is our goal to design and evaluate techniques which can adequately account for these two variables. In addition to some standard design guidelines for 3D selection techniques, this gives us the following six design guidelines:

- Allow for fast selections.
- Allow for accurate selections.
- Be easy to understand and use.
- Produce low levels of fatigue.
- Satisfy the above for sparse and dense target environments.
- Support selections for both visible and occluded targets.

While previous work will guide us in satisfying the first four design guidelines, we propose two design strategies for satisfying the last two. We now discuss these two design strategies in detail.

3.2. Increased and unambiguous activation areas

For a selection technique to be appropriate for use in an actual interface, it should support efficient selections in both sparse and dense environments. However, supporting efficient selections for both hand extension and ray casting techniques is a difficult task, due to the constraints imposed by the task. For hand extension techniques, the added third dimension imposes both a physical and visual constraint which can slow down selection times (Bowman et al., 1999; Grossman and Balakrishnan, 2004). For ray casting techniques, small angular changes in the input device can result in large movements of the ray (Liang and Green, 1994). Our strategy for reducing the effects of these constraints will be to increase the activation areas of the selection techniques. Established methods for this include: using a selection volume for hand extension techniques and emitting a cone for ray casting techniques (Liang and Green, 1994; Zhai et al., 1994).

Unfortunately, increasing the cursor's activation area means multiple targets can be captured simultaneously, introducing an undesirable ambiguity. This can be especially problematic in dense target environments. We, thus, seek 3D selection techniques which provide a mechanism for disambiguating between potential targets of interest. One strategy would be to use a dynamic activation area (Grossman and Balakrishnan, 2005). Another alternative is to provide a mechanism for disambiguating between multiple captured targets (Hinckley et al., 1994). Such techniques have a better chance of being efficient in both sparse and dense target environments.

3.3. Integrated visual enhancements for occluded targets

It is often the case in 3D environments that users need to select objects which are obscured from their viewpoint. This is generally the case when the target of interest lies behind another object in the scene and is thus occluded. If the environment is densely populated with targets, the chance of such an occlusion occurring increases. With a target being invisible to the user, it is generally impossible to select, as the user will have no visual feedback as to the location of the cursor relative to the intended target.

As discussed in Section 2.3, existing techniques for overcoming such occlusion generally require explicit and separate steps on behalf of the user. This means that the actual selection cannot be performed in a single, fluid, interaction. This is a drawback which motivates us to find other strategies.

When a user wishes to select an occluded object, the occlusion generally introduces a problem of *access*, and not *discovery* (Elmqvist and Tsigas, 2006). This is because it can be assumed that the user has a good conceptual model of the environment, knowing the general area of the intended target. As such, explicitly altering the viewing state of the entire scene may be excessive. We apply a design strategy of integrating visual enhancements into the selection technique itself. This allows occluded targets to be made visible, without the requirement of additional modes or auxiliary devices. The idea is to apply these enhancements only to targets in the vicinity of the selection cursor, almost as if the selection cursor were also controlling a magic lens (Bier et al., 1993). This strategy can be thought of as improving the *visual* feedback during the selection task. Later, in Section 6 of this paper, we will investigate the use of other feedback modalities.

4. Selection techniques

In this section, we discuss the selection techniques which will be used in our first evaluation, developed with the above mentioned design guidelines and strategies in mind. We applied these design strategies to both hand extension and ray casting metaphors, resulting in a 3D bubble cursor and an augmented depth ray.

4.1. 3D bubble cursor

We first apply our design guidelines to the hand extension metaphor. If we simply increase the activation area of such a selection technique we are left with a volume cursor, where the user must capture the intended target inside the cursor volume to select it. However, we also require the activation area to be unambiguous, and with a volume cursor, multiple targets can fall within the cursor's boundaries. To alleviate this problem, we implemented a 3D version of the bubble cursor, which dynamically resizes such that only the closest target falls within its boundaries. We render the bubble cursor as a gray semi-transparent sphere similar to the rendering of the silk cursor (Zhai et al., 1994) (Fig. 1a). When necessary, we render a second semi-transparent sphere around the captured target, such that it always appears to be fully contained by the cursor (Fig. 1b). For added visual feedback, we highlight captured targets yellow. As with the 2D implementation, we also render a crosshair inside the bubble cursor to mark its center. The user controls the location of this crosshair by positioning the input device in 3D space. We refer the reader to the original bubble cursor work for a full description of the algorithm used to calculate the cursor radius (Grossman and Balakrishnan, 2005).

Our hope is that this cursor will allow for efficient selections of targets in both sparse and dense environments. However, as per our design guidelines, we also wish for the technique to support the selection of targets which are occluded from the user's viewpoint. To overcome such occlusions, we give the bubble cursor magic lens capabilities, such that targets in its vicinity become semi-transparent. To do so, we calculate the distance between the bubble cursor and each target, measured on the 2D image viewing plane. Any target within 4 cm is rendered as semi-transparent, so that targets which may be initially occluded become visible (Fig. 2a). This allows users to hone in on an occluded goal target as they approach it, assuming they know its general location (Fig. 2b).

It is important to note that this localized transparency function is only appropriate when users know the general region of their intended target. This is often the case when users are familiar with the scene that they are interacting with, and is an assumption we make in this paper. If the user needed to search for the target, global methods such as rotating the scene or switching to a different viewing mode would be more appropriate.

4.2. Depth ray

The 3D bubble cursor results from an application of our design strategies to the hand extension metaphor. Here, we apply these guidelines to the ray casting metaphor.

Without augmentation, ray cursors can already be thought of as having increased activation areas, since they can select any target along the length of the ray. A conic ray cursor further increases this activation area. Thus, multiple targets can simultaneously be captured, and we are again required to provide a disambiguation mechanism.

As previously discussed, the depth ray has been shown to provide an effective mechanism for disambiguating between multiple intersected targets (Grossman and Balakrishnan, 2006). The user controls a depth marker, which exists along the length of ray. Moving the hand forwards or backwards will make the depth marker move in the same manner. The object intersected by the ray cursor, which is closest to the depth marker, can be selected. We render the ray as a thin red cylinder, although a conic selection area, with an apex of 1°, originating from the user's hand, is used for the selection test. As with the bubble cursor, the captured target is highlighted yellow and remaining targets intersected by the ray are highlighted green. Fig. 3 shows a screenshot of our implementation.

To allow for the selection of occluded targets, we augment the depth ray with a similar transparency function used by the bubble cursor, using the distance between the targets and the ray, measured in the 3D environment.

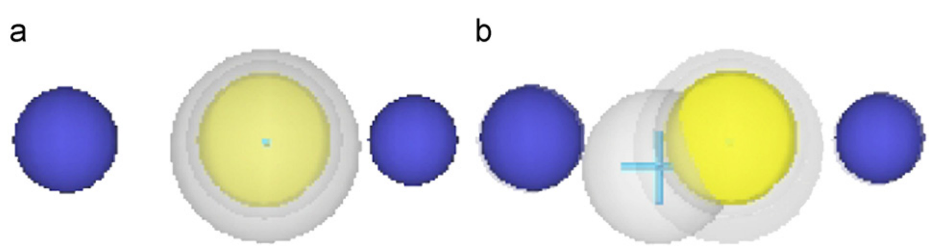

Fig. 1. (a) The 3D bubble cursor is rendered as a semi-transparent sphere which dynamically resizes such that it only captures the closest target, highlighted yellow. (b) When necessary, a second sphere is rendered around the captured target so that it always appears to be completely contained by the bubble cursor.

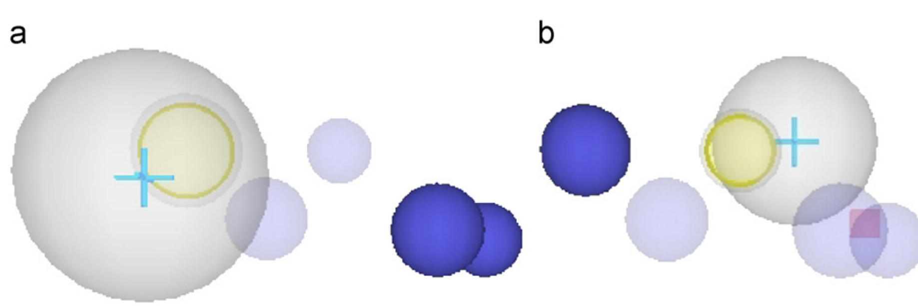

Fig. 2. (a) Targets in close proximity to the bubble cursor become semi-transparent. (b) As the cursor approaches an occluded goal target (red cube) it becomes visible and can be selected.

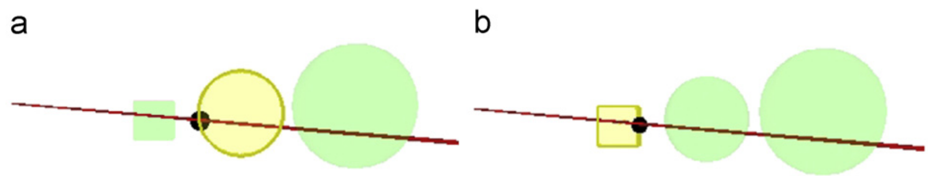

Fig. 3. (a) The depth ray selects the intersected target which is closest to the depth marker. (b) The depth marker position can be controlled by moving the hand forwards or backwards.

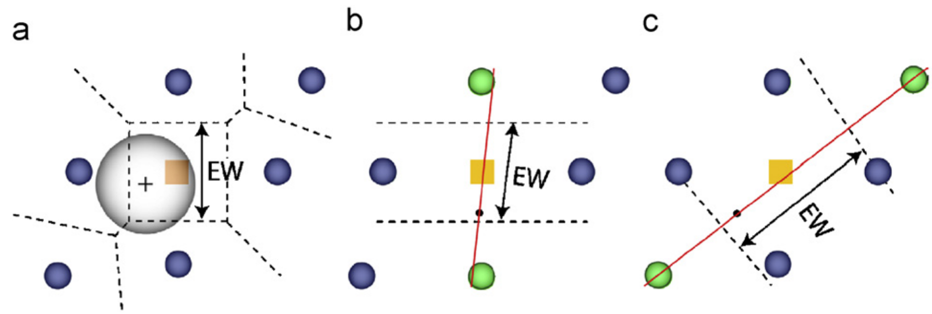

Fig. 4. (a) The bubble cursor divides the environment into 3D Voronoi regions. The effective width of each target is defined by its corresponding Voronoi region. (b) When using the depth ray, the Voronoi region is based only on the intersected targets. (c) Changing the position of the ray can change a target's effective width.

4.3. Effective widths

The effective width of a target can be defined as the size in motor space of a target's activation boundaries. It has been shown that the effective width of a target plays a larger role than its visual boundaries on selection times (Blanch et al., 2004; Grossman and Balakrishnan, 2005; Elmqvist and Fekete, 2008).While the goal of this work is not to obtain a sound theoretical model of how each technique will perform, understanding the effective widths for each technique will help us form hypotheses on the relative performances of the two techniques.

The dynamic activation area of the bubble cursor divides the 3D environment into regions, such that there is exactly one target inside each region, with that target being closest to any point inside that region. These are also known as Voronoi regions. The 3D bubble cursor increases the effective width of a target to its surrounding 3D Voronoi region. In other words, to select a target, the center of the bubble cursor only needs to be positioned inside the target's associated Voronoi region (Fig. 4a).

For the depth ray, two dimensions of the effective width depend on the angular accuracy required to intersect the target. This depends on the target size, distance, and conic apex of the ray. The third dimension of effective width depends on which surrounding targets are intersected by the ray. This establishes a Voronoi region, which when projected onto the length of the ray, defines the segment of the ray where the depth marker can select the target. It is the length of this segment which can be thought of as the effective width. In some situations, this will result in a similar effective width to when using the bubble cursor (Fig. 4b). However, if the ray is positioned to avoid certain surrounding targets, the effective width can be larger (Fig. 4c). This could potentially make the depth ray a faster technique than the bubble cursor.

5. Experiment 1: the effects of density and occlusion

We have applied our design strategies to both hand extension and ray casting metaphors, giving us the 3D bubble cursor and depth ray techniques, both augmented to allow for the selection of occluded targets. Our hope is that these techniques will allow for efficient selections in both sparse and dense target environments, and for both visible and occluded targets. Previous work tells us that the depth ray can be efficient in dense environments (Grossman and Balakrishnan, 2006), but it is unknown how exactly the density will affect performance. A better understanding exists for the effect of density on the 2D form of the bubble cursor (Grossman and Balakrishnan, 2005), but no such understanding exists for its 3D counterpart which we have introduced in this paper.

In this section, we present a formal experiment to evaluate these two techniques, where we manipulate both the environment density and goal target visibility in a 3D VR environment. One goal of the experiment is, thus, to gain an understanding of how these variables affect selections. However, another goal is to compare the relative performance of these two techniques, and also provide a comparison to a baseline technique.

In some of the experimental environments which we will use, it would be impossible to select a target with a naive implementation of the ray cursor, because of occlusions. For that reason, we use a 3D point cursor for a baseline comparison technique. The point cursor is rendered as a

3D crosshair, the same as the crosshair drawn in the center of the 3D bubble cursor, described in Section 4.1. In order to select a target, the center of the crosshair must be positioned inside of it. To allow for the selection of occluded targets, we augment the point cursor with the same transparency functionality used for the other techniques. A true baseline technique would not have this built-in functionality, and would require users to either rotate the scene or switch viewing modes to select an occluded target. However, including the transparency function minimizes hardware requirements, and provides for the fairest possible comparison.

We hypothesize that an increased environment density will slow selection times for the depth ray and bubble cursor, as an increased density will decrease effective target widths. However, the density should have little effect on a standard point cursor, as it has no effect on the motor space size of the target's activation boundaries. We also hypothesize that even with the transparency functions which our techniques use, it will still take longer to select targets which are occluded, as the visual feedback is initially reduced. However, our hope is that our techniques will reduce the overhead cost when the targets are not initially visible. Lastly, in comparing the three techniques, we hypothesize that the depth ray and bubble cursor will outperform the point cursor. The relative performance of the bubble cursor and depth ray is harder to predict, and will be an interesting result to analyze.

5.1. Apparatus

The display we used was a 21 in Hitachi CM813ETPlus monitor combined with StereoGraphics CrystalEyes 3D LCD shutter glasses for stereo viewing. Left and right eye images were provided using quad buffering at a refresh rate of 60 Hz per eye which was coordinated through the shutter glasses by an infrared transmitter positioned on top of the monitor.

For input a Polhemus Fastrak 6 DOF magnetic tracker was used, encapsulated in a handheld device. The device was equipped with a single button. The tracker was updated at 120 Hz with a precision of less than 1mm. The input device controlled each of the three cursors with an absolute 1 to 1 mapping. Participants stood during the course of the experiment, and the display was raised to be roughly shoulder level. Fig. 5 illustrates the experiment apparatus.

5.2. Participants

Eleven male and one female unpaid volunteers, ranging in age from 20 to 28, served as participants in this experiment. Participants were screened using a stereopsis test in which they had to order objects according to their depth. All participants were right handed and used their right hand to control the input device.

5.3. Procedure

A 3D static target acquisition task was used for the experiment. The scene consisted of a start target, goal target, and 45 distractor targets. The start target was rendered as a white sphere, the goal target as a red cube, and the distractor targets as blue spheres. To complete the task, participants first selected the start target, centered at the front of the display space. Once selected, this target would disappear, and the participant would then have to select the goal target. The position of the start target was static, and the goal target was positioned in a random location, such that the distance between the start and goal targets was 20 cm in 3D space. The positions of the distractor targets were random, with the constraint that they did not intersect. The radius for the distractor targets were randomly assigned values between 0.75 and 1.5 cm. Fig. 6 illustrates the experiment environment.

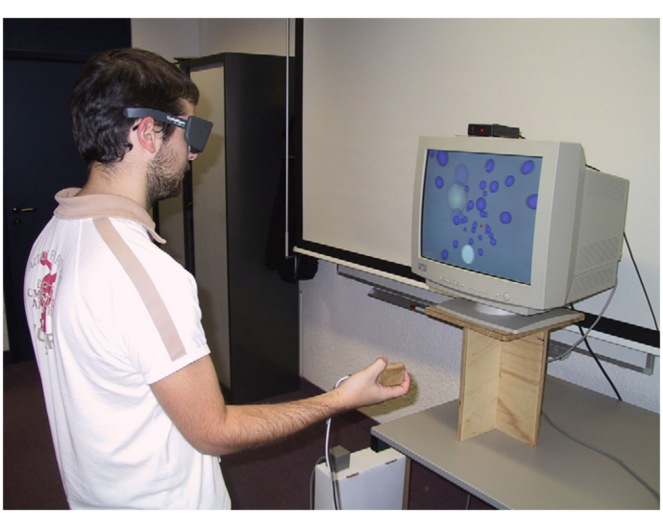

Fig. 5. The experiment apparatus.

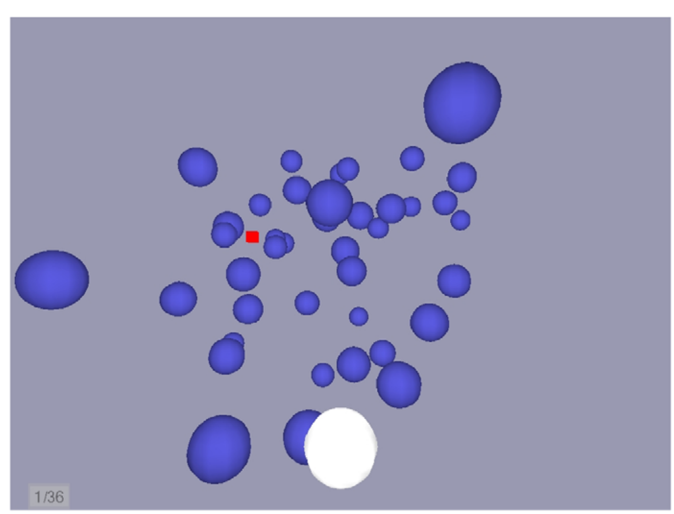

Fig. 6. The experiment environment.

Each trial started with a fade-in of the scene. Initially, only the goal target would be visible. After 500 ms the distractor targets would fade in over a duration of 2 s. This would give users an understanding of the general location of the goal target for the occluded conditions. Once all targets appeared, the user could select the start target to begin the trial. Users were instructed to select the goal target as fast as possible while minimizing their errors. Users had to successfully select the goal target before the trial could end.

For all techniques the captured target was rendered yellow with a solid opaque border. Additionally, other targets intersected by the depth ray were highlighted green.

5.4. Independent variables

The main factors we study in this experiment are the selection technique, environment density, target visibility, and target size.

For environment density, we were mostly concerned with the immediate surroundings of the goal target, as this would affect the effective width for the bubble cursor and depth ray. To control this variable, we carefully position six distractor targets around the goal target. Two distractor targets were placed along the direction of movement, defined as the vector between the start and goal targets, one before and one after the goal target. Perpendicular to this vector, targets were then placed above, below, and to either side of the goal target, forming a cube shaped Voronoi region. We controlled the size of the resulting Voronoi region by changing the distance between these six distractor targets and the goal target, measured from their closest edge. We call this distance variable density spacing, *DS* (Fig. 7).

For goal target visibility, we tested fully visible and fully occluded conditions. In both conditions, we ensured that none of the six surrounding targets were occluding the goal target. To ensure this, in some cases, we rotated the entire Voronoi region. For the visible condition, we ensured that no other distractor targets occluded the goal target. For the occluded condition, we ensured that two distractor targets partially occluded the goal target, such that together the goal target was completely occluded.

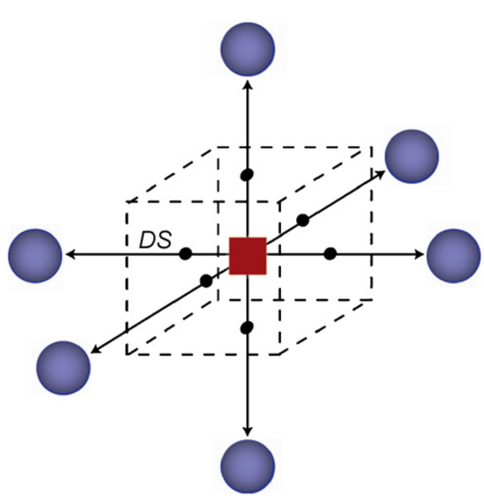

Fig. 7. Six distractor targets are carefully positioned around the goal target, creating a cube-shaped Voronoi region. The distance between these targets and goal target is the density spacing (*DS*).

5.5. Design

A repeated measures within-participant design was used. The independent variables were: cursor type, *CT* (point cursor, bubble cursor, and depth ray); density spacing, *DS* (1, 2.5, and 5 cm); visibility condition, *VC* (visible, occluded); and target size, *SIZE* (0.75 cm, 1.5 cm). A fully crossed design resulted in 36 combinations of *CT*, *DS*, *VC*, and *SIZE*.

Each participant performed the experiment in one session lasting about 70 min. The sessions were broken up by the three cursor types, with three blocks appearing for each of the techniques. Within each block, the 12 combinations of *DS*, *VC* and *SIZE* were repeated three times in a random order, for a total of 36 trials. Prior to the experiment, 36 environments were randomly generated for these trials, and were the same across all participants. The cursor ordering was fully counterbalanced across the 12 participants, with 2 participants randomly assigned to each of the 6 unique orderings. Before each cursor type, participants were given several warm-up trials to familiarize themselves with the selection technique.

5.6. Results

5.6.1. Trial completion time

In our analysis of trial completion time, we discarded trials in which errors occurred, and removed outliers that were more than three standard deviations from the group mean (1.7% of the data).

Repeated measures analysis of variance showed main effects for *CT* ($F_{2,22} = 383$), *SIZE* ($F_{1,11} = 277$), *VC* ($F_{1,11} = 330$), and *DS* ($F_{2,22} = 68.7$), all at the $p < 0.0001$ level. Average trial completion times were 4.59 s for the point cursor, 3.10 s for the bubble cursor, and 2.85 s for the depth ray. Post hoc comparisons showed that both the depth ray and bubble cursor were significantly faster than the point cursor ($p < 0.0001$), and that the depth ray was significantly faster than the bubble cursor ($p < 0.005$).

We also found that *CT* had significant interaction effects with *SIZE* ($F_{2,22} = 12.8$, $p < 0.0001$), *VC* ($F_{2,22} = 118$, $p < 0.0001$), and *DS* ($F_{4,44} = 3.97$, $p < 0.005$), showing that each of these independent variables affected the three techniques differently. We now provide a discussion on each of these observed interaction effects.

The interaction between *CT* and *DS* is illustrated in Fig. 8. The general trend, as expected, is that the selection times are reduced with increased density spacing. There are, however, two interesting results which are illustrated.

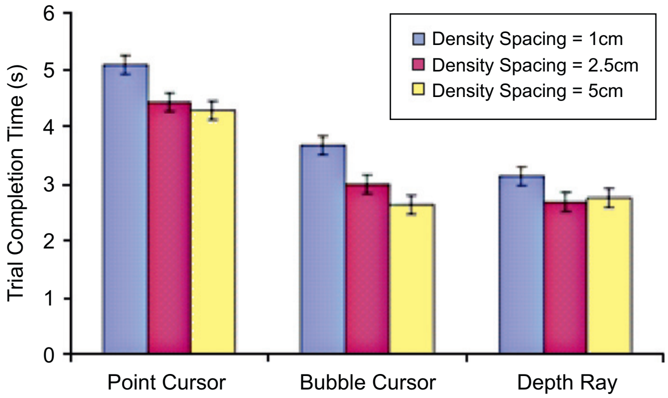

Fig. 8. Technique completion times by density.

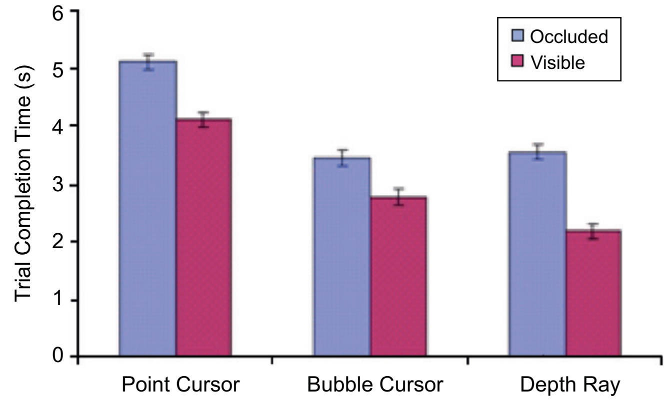

Fig. 9. Technique completion times by visibility condition.

First, the density had a significant effect on movement times for the point cursor ($p<0.0001$). This is somewhat surprising since it should only be the target size which constrains a selection with the point cursor. A possible explanation is inadequate visual feedback. In the dense conditions, it may have been more difficult to discern when the goal target had been highlighted.

Second, it is interesting to note that for the depth ray, increasing the density spacing from 2.5 to 5 cm did not significantly reduce movement times. In fact, the movement times actually increased, although the effect was not significant. In contrast, the bubble cursor completion times decreased significantly for each increase in density spacing ($p<0.05$). While the data for the depth ray may at first seem counterintuitive, we must recall that the effective width of the goal target when using the depth ray is not completely determined by the density spacing value. It also depends on the angle of approach that the ray takes on (Fig. 4), so it may be the case that users were not taking on optimal approach angles for the condition $DS = 5$ cm. This could have been for many reasons, one of which being that the randomly generated environments were different for each density value.

Fig. 9 illustrates the interaction between *CT* and *VC*. As can be seen, the occluded condition increases trial completion times for all three cursors ($p<0.0001$), with an average increase in completion times of 1.01 s. From this figure, we can also see that the depth ray was the fastest technique overall because of its superiority in the visible condition. Post hoc multiple means comparison shows that in this condition, the depth ray is significantly faster than the bubble cursor ($p<0.0001$), whereas in the occluded condition, there is no difference between the two techniques. It may be the case that the overhead cost introduced by the occluded condition outweighs the differences between the two techniques, so their differences are not apparent.

Finally, the interaction effect between *CT* and *SIZE* is illustrated in Fig. 10. As can be seen, the size had its most dramatic impact on the point cursor, which is the only technique constrained by the dimensions of the actual target. However, the size also had a significant effect on the bubble cursor ($p<0.0001$), with completion times of 3.4 s for $SIZE = 0.75$ cm and 2.8 s for $SIZE = 1.5$ cm. This is somewhat surprising, since selections with this technique are constrained by the proximity of the surrounding targets, and not the actual target size. This indicates that users were being drawn towards the visible borders of the targets, not taking full advantage of the target's effective width. This result is consistent with observations of the 2D bubble cursor, which showed a similar result, although not as strong. The effect may be increased in our 3D task because the visual feedback is not as clear, as multiple objects could be layered over the cursor.

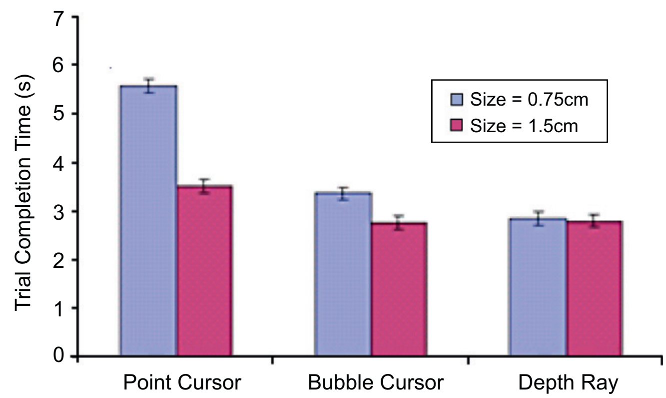

Fig. 10. Technique completion times by target size.

In Grossman's original bubble cursor study (Grossman and Balakrishnan, 2005), it was shown that the size of the target will have more effect on selection times when the surrounding targets are closer. To determine if this was the cause for the effect we were observing, we looked at the interaction between *SIZE* and *DS* for the bubble cursor condition. Indeed, the interaction effect was significant ($F_{2,22} = 17, p<0.0001$), illustrated in Fig. 11. It can be seen that the effect of size increased when surrounding targets were closer to the goal target. Post hoc comparisons show that size had a significant effect for $DS = 1$ cm and 2.5 cm ($p<0.0001$), but not for $DS = 5$ cm.

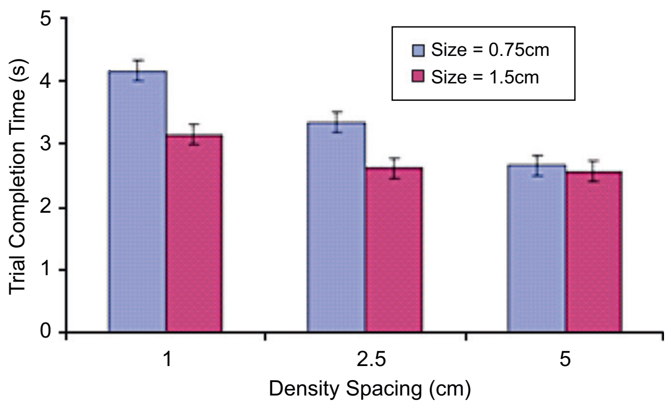

Fig. 11. Bubble cursor times by size and density spacing.

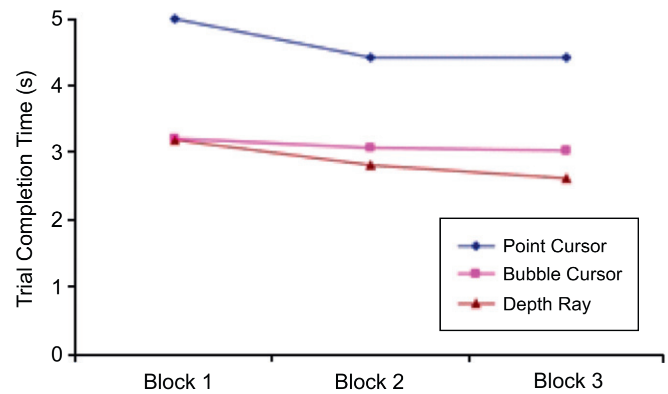

Fig. 12. Learning effects for each of the techniques.

As for the depth ray, target size had no significant effect, indicating the visual feedback of when targets were captured may have been better than for the bubble cursor, reducing unnecessary movements towards the boundaries of the goal target.

5.6.2. Learning

A learning effect was observed in our experiment, with the block number significantly affecting trial completion times ($F_{2,22} = 23.8$, $p < 0.0001$). Performance improved over time, with average completion times of 3.79, 3.43, and 3.33 s for blocks 1, 2, and 3, respectively. Post hoc comparisons shows that Block 1 was significantly slower than Blocks 2 and 3 ($p < 0.0001$), but Blocks 2 and 3 were not significantly different.

There was also a significant interaction between the block number and *CT* ($F_{4,44} = 2.54$, $p < 0.05$). The effects, illustrated in Fig. 12, indicate that there is little overhead cost involved with learning the techniques. It can be seen that the most learning occurs with the point cursor. It is also interesting to note that the bubble cursor and depth ray are almost equivalent in the first block, and differ significant only in the last block ($p < 0.05$). This may indicate that users learned to take optimal approaches to the goal target using the depth ray, to maximize the effective width.

5.6.3. Input device footprint

The input device footprint is defined as the length of the total path which the input device traveled to complete each trial. The cursor type had a significant effect on the input device footprint ($F_{2,22} = 210$, $p < 0.0001$). The footprints were 31.2 cm for the point cursor, 26.9 cm for the bubble cursor, and 22.4 cm for the depth ray, all significantly different ($p < 0.0001$). The depth ray likely had the lowest value because, for this technique, users are not required to move their hand towards a specific 3D location.

5.6.4. Error rates

For this task, errors were defined as trials in which users selected the wrong target before selecting the goal target. The overall error rate for the experiment was 2%. For the point cursor no errors were reported. This shows that users could use the visual feedback provided to tell when the goal target was highlighted. The bubble cursor had an error rate of 2.2% and the depth ray 3.4%, for which there was no statistical difference.

5.6.5. Subjective feedback

Overall, participants found the bubble cursor and depth ray easy to use, preferring these techniques over the point cursor. Participants who preferred the depth ray noted that less movement was required, and that they could aim from any angle. All 12 participants responded that they liked the transparency function.

6. Multimodal feedback

The results from our first experiment indicate that while the provided visual feedback works adequately, it does not completely alleviate the difficulties imposed by selections within dense and occluded target environments. For example, users did not take full advantage of the effective target width when using the bubble cursor, and relied on the visual boundaries of the target. Furthermore, with all techniques, there was about a 1 s overhead cost when selecting targets which were occluded from the user's viewpoint.

These shortcomings motivate us to investigate other forms of feedback, to compliment the visual feedback which we provided, so that users will have a better overall understanding during the selection task. In a follow-up experiment which we will describe in the next section, we investigate if other forms of feedback, specifically audio and force, could further aid users during selections within dense and occluded target environments.

We will only investigate the addition of these forms of feedback to the ray casting and bubble cursor selection techniques, and no longer consider the point cursor, because it was found to be less efficient, and it is unlikely that multimodal feedback would provide it with enough of a performance gain to surpass the other techniques.

Furthermore, multimodal feedback techniques for the traditional point cursor have already been investigated in both 2D and 3D environments (Cockburn and Brewster, 2005; Vanacken et al., 2006). In contrast, the use of multimodal feedback for selection techniques with dynamic activation areas, such as the bubble cursor and depth ray, to our knowledge, has never been thoroughly investigated.

Because the selection techniques we are using have dynamic activation areas, the main purpose of the additional feedback is to make the user more aware of the targets' activation boundaries. This means that the feedback needs to be activated when the user switches between targets. To prevent the user from being overwhelmed with feedback during the initial ballistic movement of selection (Meyer et al., 1988; Rosenbaum, 1991), we disable the feedback during fast movements. We use both a velocity and acceleration threshold (3.0 cm/s and $1.0\,cm/s^2$, respectively). When either threshold is crossed, we disable the additional feedback (Oakley et al., 2001). We now discuss the additional feedback which we will apply to our selection techniques.

6.1. Force feedback

It has been shown that attraction forces, commonly used when haptic feedback is used during selection, can be detrimental, especially in dense target environments (Hwang et al., 2003; Keuning, 2003; Ahlström, 2005; Ahlström et al., 2006). Since our motivation is to find techniques which work well in such environments, we will diverge from these magnetisms techniques. Instead, we propose a short and subtle haptic response during selection. It is best described as a quick bump in the movement with a low force to make sure that the user's movement is not disturbed. For the bubble cursor, the direction of the force feedback is parallel to a Voronoi edge with a duration of 25 ms and maximum strength of 1 N with a sine wave profile. For the depth ray, a bump, with the same properties, is felt in the direction perpendicular to the ray. This bump is felt any time the movement of the depth marker causes a new target to be captured.

6.2. Audio feedback

The audio feedback we use is the earcons technique which has been traditionally used (Akamatsu et al., 1995; Cockburn and Brewster, 2005). Whenever a new target is hit, an earcon is played for a short duration. We do not play the earcon of a long duration (Cockburn and Brewster, 2005) or continuously while a target is being hit (Akamatsu et al., 1995), because with our techniques of interest a target is always selected, which would result in continuous, and thus meaningless, audio feedback. Any time a new target is captured, an audio earcon (C on celesta for 0.2 s) was sounded.

6.3. Visual feedback

In addition to the force and audio feedback, we also changed some properties of the visual feedback provided during the experiment. These modifications were made based on observations and user feedback from Experiment 1. In the first experiment, we used a discrete transparency function, which was either enabled or disabled for each target based on whether that target was within a threshold distance of the cursor. We modified this to use a continuous function instead. With this function the percentage of transparency depends on the relative distance, which prevents what could sometimes be sudden and distracting changes in target alphas. We also changed the visual appearance of the bubble cursor. We decided to not render the bubble at all, and only render the crosshair representing its center. We hoped this would make it easier to see which target was highlighted and thus captured.

7. The haptic lock ray

We have augmented our earlier designed techniques, the bubble cursor and depth ray, with new forms of multimodal feedback. During our initial tests with the depth ray we found that too many parameters were changing simultaneously (such as the trajectory of the ray, location of the depth marker, proximity to targets, etc) for the multimodal feedback to seem intuitive to users. The lock ray (Grossman and Balakrishnan, 2006), a technique closely related to the depth ray, splits the selection task into explicit selection and disambiguation phases, and would thus be potentially more appropriate with the presence of multimodal feedback. With this technique, the depth marker only becomes active once the user presses the input device button. At this point the user can control the depth marker, as they would with the depth ray, but the position of the ray itself is locked. The user confirms the selection by releasing the button. We slightly modify the originally developed lock ray by initially placing the depth marker at the midpoint of all intersected targets, to minimize its expected travel distance. Furthermore, we render the depth marker when it is inactive, so the user will know exactly where it is before they begin to use it. When inactive the marker is grayed out. The disambiguation phase, when the user is controlling the location of the depth marker and the position of the ray is locked, is ideal for multimodal feedback design. We used the same audio feedback as with the other selection techniques, except that an additional earcon (*F* on piano for 0.2 s) signaled that the disambiguation phase had started. During the disambiguation phase, the haptic input device was constrained to move only along the vector of the ray. In addition the user felt similar bumps which were used with the depth ray, every time a new target was selected as a result of moving the depth marker. We slightly increased the force to 1.5 N since there would likely be less targets to pass over using this technique, since the user would have already indicated the main areas of interest. This greater bump force created a subtle "pop-through" effect (Smyth and Kirkpatrick, 2006). If the depth marker reached the first or last intersected target, a force wall would prevent any further movements of the depth marker, such that potential error corrections during the disambiguation phase would be faster.

8. Experiment 2: the effects of multimodal feedback

In this section, we present a second experiment to determine if the additional multimodal feedback which we have introduced can improve users' ability to select targets in dense and occluded 3D environments. We will compare and evaluate the bubble cursor, depth ray, and lock ray, both in the presence and absence of the new multimodal feedback. We omit a comparison with our previous baseline factor the point cursor because it proved to be much slower than both the bubble cursor and depth ray. The procedure of the experiment was identical to Experiment 1.

8.1. Apparatus

The apparatus differed from Experiment 1 as this experiment was carried out at a later time, and a haptic input device was required. The display was a 2.4 m × 1.8 m polarization projection screen with passive stereo using two DLP projectors. For input a PHANToM premium 1.0 was used with a stylus for 6 DOF input and 3 DOF force feedback. The stylus is equipped with a single button. The force update rate was 1000 Hz and the tracker update rate was 120 Hz. The input device controlled each of the three cursors with an absolute 1 to 0.35 mapping. Participants were seated during the experiment 3 m from the display and wore normal headphones at all times. Fig. 13 illustrates the experiment apparatus.

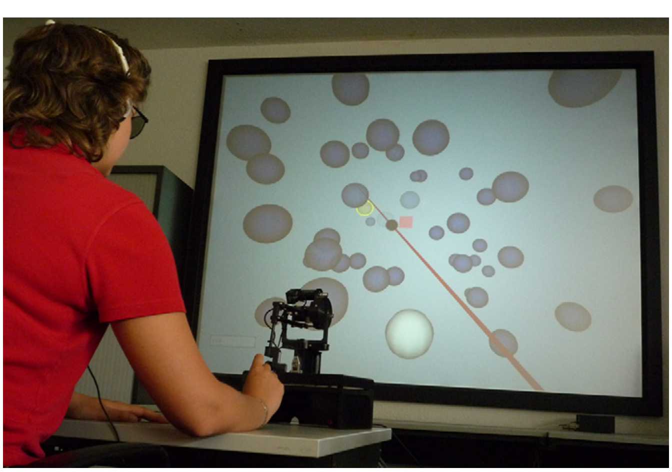

Fig. 13. The experiment apparatus (stereo is disabled for illustrative purposes).

8.2. Participants

Eleven male and one female unpaid volunteers, ranging in age from 21 to 30, served as participants in this experiment. Participants had not taken part in the previous experiment and were screened using a stereopsis test in which they had to order objects according to their depth. All participants were right handed and used their right hand to control the input device.

8.3. Design

A repeated measures within-participant design was used. The independent variables were: cursor type, *CT* (bubble cursor, depth ray, and lock ray); multimodal feedback, *MM* (on, off); density spacing, *DS* (1 and 5 cm); visibility condition, *VC* (visible, occluded); and target size, *SIZE* (0.75 and 1.5 cm). A fully crossed design resulted in 48 combinations of *CT*, *MM*, *DS*, *VC*, and *SIZE*. The *DS*, *VC*, and *SIZE* variables behaved exactly the same as in Experiment 1. The multimodal feedback variable indicated the presence of both force and audio feedback. The same visual feedback was used in all conditions.

Each participant performed the experiment in one session lasting about 70 min. The session was broken up by the three cursor types. Trials for each cursor were further split between the two values of *MM*. For each *CT* and *MM* combination, three blocks of trials were performed. Within each block, the six combinations of *DS*, *VC* and *SIZE* were repeated three times in a random order, for a total of 18 trials per block. In our previous experiment, 36 environments were randomly generated for these trials, and were the same across all participants. We reused 18 of those 36 environments which fit the conditions of this experiment. The cursor ordering was fully counter-balanced across the 12 participants, with two participants randomly assigned to each of the six unique orderings. The multimodal feedback was first off and then on for each cursor for the first group of six participants, and this order was swapped ($MM =$ on, $MM =$ off) for the second group of six participants. Before each cursor, participants were given several warm-up trials to familiarize themselves with the selection technique.

8.4. Results

8.4.1. Trial completion time

In our analysis of trial completion time, we discarded trials in which errors occurred, and removed outliers that were more than three standard deviations from the group mean (1.6% of the data). Repeated measures analysis of variance showed main effects for *CT* ($F_{2,22} = 9.08$, $p < 0.001$), *SIZE* ($F_{1,11} = 215$, $p < 0.0001$), *VC* ($F_{1,11} = 167$, $p < 0.0001$), and *DS* ($F_{1,11} = 126$, $p < 0.0001$). No effect was found for the *MM* condition ($F_{1,11} = 0.018$, $p = .896$). Average trial completion times were 2.34 s for the bubble cursor, 2.44 s for the depth ray, and 2.76 s for the lock ray.

Post hoc comparisons showed that both the bubble cursor and depth ray were significantly faster than the lock ray ($p<0.005$) and the difference between the bubble cursor and the depth ray was not significant ($p = 0.382$). Despite this non-significance, the fact that bubble cursor times were now faster than the depth ray overall, in contrast to Experiment 1, where they were significantly slower, indicates that the adjusted visual feedback, used for the bubble cursor in this experiment, was effective. The finding that the lock ray was slower than the depth ray is consistent with the prior literature (Grossman and Balakrishnan, 2006), and indicates that the addition of multimodal feedback is not enough to overcome their differences.

Another important result is that the multimodal feedback did not have any significant effect or interaction with the technique (*CT*), as illustrated in Fig. 14. This was somewhat disappointing; our hope was that the multimodal feedback would further improve the users' awareness during the selection task. But as a positive note, this indicates that the visual feedback provided with the techniques, alone, is sufficient. The only effect of interest related to the multimodal feedback was that *MM* had a weak but significant interaction with *DS* ($F_{1,11} = 6.39$, $p = 0.028$). Fig. 15 shows that this effect is barely apparent, but it does indicate that the multimodal feedback could be more useful in dense target environments, when it becomes harder to discern the targets visually. We further investigated the effect of *MM* by looking at reaction times, as this is a metric that multimodal feedback has been observed to have an effect on in past studies (Akamatsu et al., 1995). It is defined as the time taken to select an object with a button press once it has been captured with the cursor. However, a repeated measures analysis of variance showed no main effect of *MM* on reaction time either ($F_{1,11} = 0.430$, $p = 0.525$).

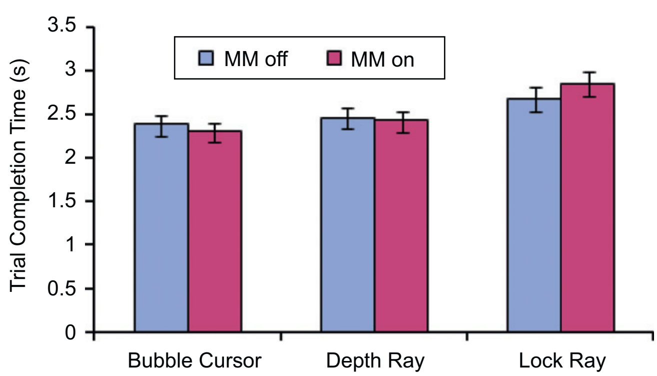

Fig. 14. Technique completion times by multimodal feedback.

We now discuss how the other variables—target size, visibility, and density—influenced the techniques. We found that *CT* had significant interaction effects with each of these variables: *SIZE* ($F_{2,22} = 12.8$, $p<0.0001$), *VC* ($F_{2,22} = 118$, $p<0.0001$), and *DS* ($F_{2,22} = 16.6$, $p<0.0001$).

The interaction between *CT* and *DS* is illustrated in Fig. 16. The general trend, as expected, is that the selection times are reduced with increased density spacing. It is interesting that the bubble cursor is significantly faster than the depth ray and lock ray for density spacing of 5 cm ($p<0.008$ with bonferroni corrections), while for density spacing of 1 cm there is no significant difference between any of the three cursors. This is in contrast to Experiment 1, where the bubble cursor was slower than the depth ray at $DS = 1$ cm, and similar at $DS = 5$ cm. This again indicates that the new visual feedback for the bubble cursor was effective.

Next, we discuss the interaction between *CT* and *VC*. In Fig. 17 we can see that the occluded condition is slower for each cursor. Post hoc comparisons with bonferroni correction show that it is significant for all cursors ($p<0.001$). Most notable is that the bubble cursor seems much less affected by the visibility condition than the ray casting techniques. This shows that the 3D bubble cursor, with the new visual feedback which we introduced, is an effective selection technique when targets may be visible or occluded. We cannot be certain as to why the ray casting techniques were less efficient in the occluded conditions. One possible explanation is that the continuous transparency function was less intuitive, or more difficult to take advantage of, when the metric was based on a distance to a

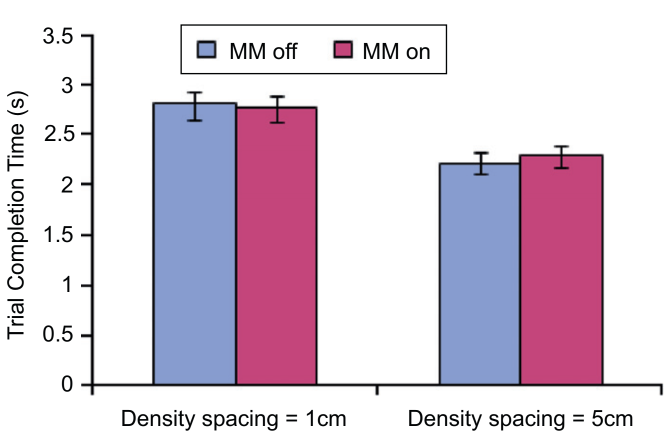

Fig. 15. Completion times by density and multimodal feedback.

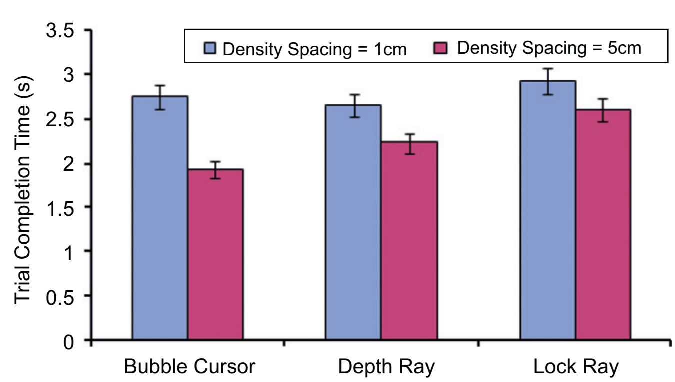

Fig. 16. Technique completion times by density.

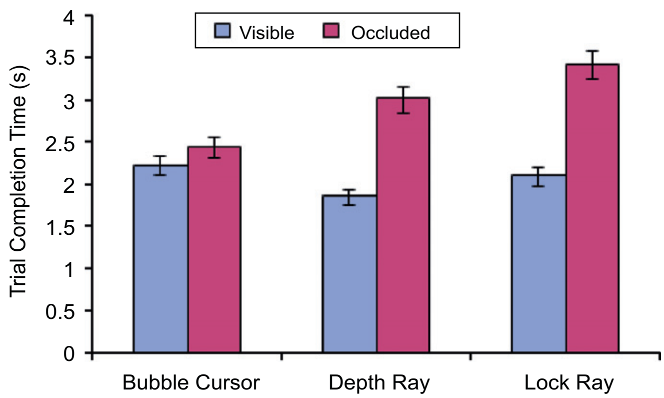

Fig. 17. Technique completion times by visibility.

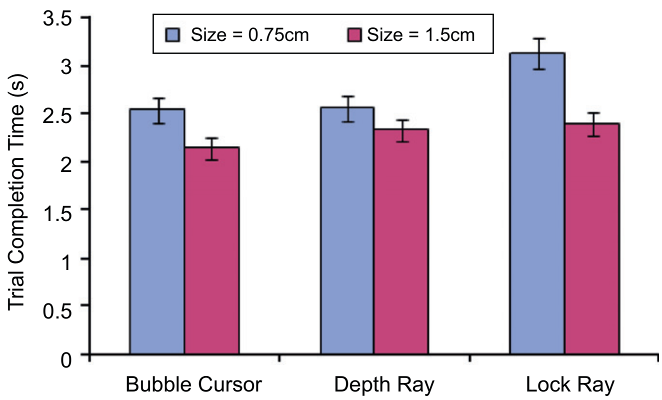

Fig. 18. Technique completion times by target size.

ray, rather than a distance to the center of the bubble cursor.

Finally, Fig. 18 shows the interaction effect between *CT* and *SIZE*. It is interesting to note that size has an effect on each of the three techniques ($p<0.0001$), even though size plays no role in the effective width, or motor space activation boundaries. Thus, as in Experiment 1, we are seeing that users are not taking full advantage of the increased activation boundaries created by the techniques. As in Experiment 1, users were more reliant on the visual boundaries of the targets when the targets were closer together. Fig. 19 illustrates this effect with the bubble cursor, which is almost identical to Fig. 11 from Experiment 1. Our hope was that the addition of the multimodal feedback would compensate for this effect, improving the users understanding of the motor space activation boundaries, but this was not the case. This indicates that in dense and occluded environments, to optimize the user's understanding of the activation boundaries of a desired target, feedback techniques—visual, haptic and audio—even combined are not enough.

The $CT \times SIZE$ interaction effect seems to be a result of the higher times for the Lock Ray for $SIZE = 0.75$ cm. This may be due to the fact that if users do rely on the visual boundaries of the target, then they have to aim at these smaller targets twice with the lock ray, once with the ray in the selection phase, and once with the depth marker

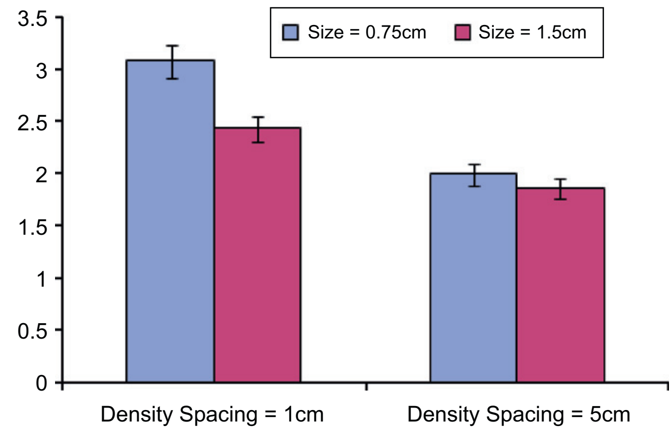

Fig. 19. Bubble cursor times by density and target size.

in the disambiguation phase. The original study of the lock ray (Grossman and Balakrishnan, 2006) did not vary size, so this is an interesting result.

8.4.2. Learning

A learning effect was observed in our experiment, with the block number significantly affecting trial completion times ($F_{2,22} = 16.2$, $p<0.0001$). Performance improved over time, with average completion times of 2.75, 2.59, and 2.46 s for blocks 1, 2, and 3, respectively. Post hoc comparisons shows that Block 1 was significantly slower than Blocks 2 ($p<0.05$) and 3 ($p<0.001$), and Block 2 was also significantly slower than Block 3 ($p<0.0001$). There was no significant effect between block number and *CT*, or block number and *MM*, indicating the techniques were just as easy to learn with or without the multimodal feedback.

8.4.3. Input device footprint

Because selection with a 6 DOF input device requires a user to hold the device in mid-air, it is important that the selection techniques minimize arm fatigue. With the device used in this experiment, a PHANToM, the user rests his arm using his elbow or under arm depending on what is most comfortable for him/her. As in Experiment 1, we can quantify potential fatigue by measuring the input device footprint, defined as the length of the total path which the input device traveled to complete each trial. The cursor type had a significant effect on the input device footprint ($F_{2,22} = 45.6$, $p<0.0001$). The footprints were 21.2 cm for the bubble cursor, 15.8 cm for the depth ray, and 8.7 cm for the lock ray, all significantly different ($p<0.005$). The lock ray may have had the lowest value because when the users completed the selection phase, the depth marker was positioned at an optimal location, in the middle of all targets being intersected. Furthermore, movements during the disambiguation phase of the lock ray were constrained to the direction of ray, which may have further reduced the footprint. As with the previous experiment, the depth ray had a lower footprint than the bubble cursor, probably because users did not need to move towards a specific 3D location.

An interesting observation was that the $CT \times VC$ interaction effect was significant ($F_{2,22} = 12.75$, $p < 0.0001$). For both the depth ray and lock ray, footprints increased significantly for $VC =$ occluded ($p < 0.001$). However, VC had no effect at all on the footprints for the bubble cursor. This further supports our understanding of why the bubble cursor was a better technique for occluded targets. The constant device footprint indicates that users were better at using the transparency function with the bubble cursor to locate the occluded targets. The increased footprint for the ray casting techniques indicates that users moved the cursors more to use the transparency function to locate the target.

8.4.4. Error rates

As in Experiment 1, errors were defined as trials in which users selected the wrong target before selecting the goal target. The overall error rate for the experiment was 2.89%. The CT did have a significant effect on error rate ($\chi^2(2) = 6.51$, $p < 0.05$), however, the overall error rates were comparable: the bubble cursor had an error rate of 1.9%, the depth ray 3.12% and the lock ray 3.64%. Post hoc analysis showed a significant difference between the bubble cursor and both the depth ray and lock ray ($p < 0.05$). The MM condition did not have a significant effect on the error rate ($\chi^2(1) = 2.99$, $p = 0.22$).

8.4.5. Subjective feedback

Users filled out a short questionnaire after each cursor had been tested. An overview of the results can be found in Fig. 20, with the results represented by median values. We see overall the results were stronger for the bubble cursor and depth ray, which again indicates that even with multimodal feedback, the lock ray is not a preferable technique. Users found that the bubble cursor and depth ray were easy to learn, and allowed for fast selections. It is also interesting to note that the results indicate that users did feel that the multimodal feedback was helpful and preferred, even though our quantitative analysis revealed that the multimodal feedback provided no advantage.

In post-experiment discussions, users reported that the MM condition helped them to better understand when a new target was selected and that it was nice to have the extra feedback. Some also commented that they were only relying on one of haptic or audio feedback, and that having both did not seem beneficial. One user did feel that neither form of feedback was at all helpful. For both the depth ray and bubble cursor users gave general remarks that the forces were sometimes disturbing during general movement. For the lock ray users found the force feedback intuitive and understood it well, but some users commented that the boundary force to switch between objects was too strong.

9. Comparison between experiments

In this section, we compare the results of the two experiments which we have presented. In order to perform the analysis, we have removed certain data from both experiments, such that the conditions which are compared are equivalent. Because the point cursor was only used in the first experiment, and the lock ray was only used in the second, we have removed the data from both techniques. Furthermore, the level of $DS = 2.5$ cm, which only appeared in the first experiment, and the condition $MM =$ on, included only in the second experiment, are both removed. These changes result in a repeated measures analysis of variance with a between-participant factor EXP, representing the experiment number.

A repeated measures analysis of variance showed a main effect for EXP ($F_{1,22} = 15.8$, $p < 0.001$). Comparing the average completion times showed that the first experiment (3.1 s) was slower than the second experiment (2.5 s). We believe the second experiment was faster because of the different hardware setup. Most notably, the control gain was altered, and a larger display was used. The faster times

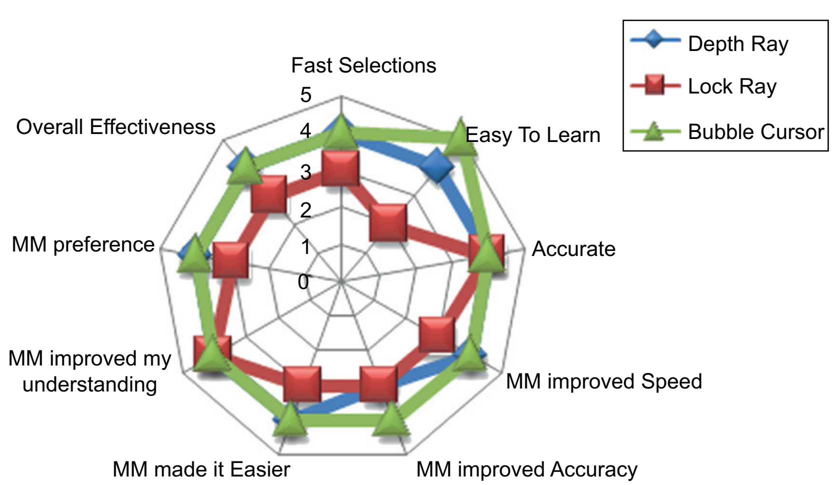

Fig. 20. Questionnaire scores from the subjective feedback.

may have also been a result of the improved visual feedback used for the bubble cursor.

Aside from this overall difference in completion times, no significant two-way interaction effect across both experiments was found. This analysis shows that while our change in setup had an overall effect across all conditions, the results which we have obtained in both experiments are consistent with one another.

10. Discussion

The results of our studies have some important implications to user interfaces for VR environments. We have discussed the design of two beneficial techniques for object selection. The poor performance of the 3D point cursor in Experiment 1 validates our design of new techniques, as well as our discussed design guidelines. While the depth ray was an existing technique for volumetric displays, we have provided its first VR implementation. Similarly we have provided the first implementation of a 3D bubble cursor. Furthermore, these techniques are both original in that they are augmented to allow for the selection of occluded targets.

In the first experiment we found that the depth ray seemed to perform best overall, and that users were not taking full advantage of the increased activation areas created by the bubble cursor. We felt this may have been due to the difficulties in providing adequate visual feedback during selection in dense and occluded targets, which led us to some changes in the visual feedback which the techniques provided. For the bubble cursor we removed the rendering of the outer bubble, and for all techniques we switched to a continuous, rather than discrete, on or off, transparency function. This seemed to make a difference in our second experiment. Overall, we felt the bubble cursor performed best in this experiment. The average completion times were not significantly different from the depth ray, but the technique was much less affected by the occlusion condition, and the error rates were lower. However, a benefit of the depth ray, found in both experiments, was that it lowered the input device footprint, which could minimize arm fatigue over extended usage.

Unfortunately, our added multimodal feedback did not provide any observable significant advantages. Our hope was that this feedback could mitigate the difficulties associated with only providing visual feedback. Mainly, users seemed to rely on the visual bounds of targets, not taking full advantage of the increased motor activation boundaries. The result that the multimodal feedback did not provide an advantage, while unfortunate, is consistent with prior art, which shows that effects multimodal feedback during selections can be minimal (Akamatsu et al., 1995; Wall et al., 2002; Cockburn and Brewster, 2005). The problem may be that providing feedback, regardless of the form, only helps the user know when they have captured a target, and does not help them plan how to capture it. Or, in other words, the users may not understand the motor activation boundaries until the target has already been captured. It would, thus, be interesting in the future to consider *feedforward* techniques, which give the user an understanding of the activation boundaries before the selection task even begins. For example, the starburst technique renders each target's surrounding activation area, which are modified Voronoi regions (Baudisch et al., 2008). It would be interesting, but potentially difficult, to adapt this strategy to 3D, without making the environment too visually cluttered. In general, it would be worthwhile to continue to investigate improvements to the visuals provided during the selection ask, as our results show that this plays a significant role in how well the users will be able to perform the task.

It is not surprising that our results showed that for all techniques, occluded targets took longer to select. However, it is important to note that our transparency function did enable users to minimize the overhead cost for such tasks to just one second. If a more traditional approach for making the target visible were used, such as using a hotkey to switch viewing modes or rotating the scene, we would expect to see a similar, if not greater, overhead cost, with the added drawback of extra buttons or input devices being required. This transparency function could also be used for higher level tasks such as the exploration of dense and complex environments.

While the addition of multimodal feedback had no significant influence, users did generally like the extra feedback. If the user's task did not otherwise require a haptic device, it would probably not be worth introducing this device just to provide the feedback, given its limited impact. However, the audio feedback could be incorporated for user's who preferred it, without any overhead costs.

11. Conclusion

We have presented an in-depth evaluation of techniques which support selections in dense target environments, and of targets which are fully obscured from the user's viewpoint. In particular, we found two techniques, the depth ray and 3D bubble cursor, which outperformed a baseline point cursor in our experimental task, providing faster selection times and lower device footprints. In a second experiment, we found that the addition of multimodal feedback to these techniques had negligent effects, although it was received well by users. The results of the second experiment also showed that our iteration on the visual feedback provided during the selections was helpful. These results indicate that interface designers should play close attention to the visual stimuli provided in their designs, and if done properly, should not have to rely on specialized hardware to provide other forms of feedback. In summary, we believe that the contributions of this work will be valuable for future designers of interactive 3D applications.

Acknowledgments

We thank the experiment participants and Tom De Weyer for his assistance in our implementation. Part of the research at EDM is funded by ERDF (European Regional Development Fund), the Flemish Government and the Flemish Interdisciplinary Institute for Broadband Technology (IBBT). We would also like to thank Doug Bowman, Bernd Froehlich, Yoshifumi Kitamura, Wolfgang Stuerzlinger, and Sabine Coquillart, editors of the International Journal of Human Computer Studies special issue on 3DUIs, for inviting this submission, and the journal reviewers, for providing insightful comments and suggestions throughout the reviewing process.

References

Ahlström, D., 2005. Modeling and improving selection in cascading pull-down menus using Fitts' law, the steering law and force fields. In: ACM CHI Conference on Human Factors in Computing Systems, Portland, Oregon, USA. ACM Press, New York, pp. 61–70.

Ahlström, D., Hitz, M., Leitner, G., 2006. An evaluation of sticky and force enhanced targets in multitarget situations. In: Proceedings of the Fourth Nordic Conference on Human–Computer Interaction: Changing Roles, Oslo, Norway. ACM, New York, pp. 58–67.

Akamatsu, M., MacKenzie, I.S., Hasbrouc, T., 1995. A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. Ergonomics 38, 816–827.

Arsenault, R., Ware, C., 2000. Eye–hand co-ordination with force feedback. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, The Hague, The Netherlands. ACM, New York, pp. 408–414.

Baudisch, P., Zotov, A., Cutrell, E., Hinckley, K., 2008. Starburst: a target expansion technique that can handle targets organized in clusters. In: Proceedings of the Working Conference on Advanced Visual Interfaces, Napoli, Italy, pp. 129–137.

Bier, E., Stone, M., Pier, K., Buxton, W., DeRose, T., 1993. Toolglass and magic lenses: the see-through interface. In: ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques. ACM, New York, NY, pp. 73–80.

Blanch, R., Guiard, Y., Beaudouin-Lafon, M., 2004. Semantic pointing: improving target acquisition with control-display ratio adaptation. In: ACM CHI Conference on Human Factors in Computing Systems, pp. 519–525.

Bolt, R., 1980. Put-that-there: voice and gesture at the graphics interface. In: ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques, vol. 14, pp. 262–270.

Bowman, D.A., Johnson, D.B., Hodges, L.F., 1999. Testbed evaluation of virtual environment interaction. In: ACM VRST Symposium on Virtual Reality Software and Technology, ACM, New York, pp. 26–33.

Bowman, D.A., Kruijff, E., LaViola, J.J., Poupyrev, I., 2004. 3D User Interfaces: Theory and Practice. Addison Wesley Longman Publishing Co., Inc., Reading, MA.

Brewster, S., 1998a. The design of sonically-enhanced widgets. Interacting with Computers 11, 211–235.

Brewster, S., 1998b. Using earcons to improve the usability of a graphics package. In: Proceedings of HCI on People and Computers XIII. Springer, Berlin, pp. 287–302.

Buxton, W., Fitzmaurice, G.W., 1998. HMD's, caves, & chameleon: a human-centric analysis of interaction in virtual space. Computer Graphics, The SIGGRAPH Quarterly 32, 64–68.

Carpendale, S., Cowperthwaite, D., Fracchia, D., 1997. Extending distortion viewing from 2D to 3D. IEEE Computer Graphics and Applications: Special Issue on Information Visualization 17, 42–51.

Cockburn, A., Brewster, S., 2005. Multimodal feedback for the acquisition of small targets. Ergonomics 48, 1129–1150.

De Boeck, J., De Weyer, T., Raymaekers, C., Coninx, K., 2006. Using the non-dominant hand for selection in 3D. In: IEEE Symposium on 3D User Interfaces, pp. 53–58.

Elmqvist, N., 2005. Balloon Probe: reducing occlusion in 3D using interactive space distortion. In: ACM VRST Symposium on Virtual Reality Software and Technology, Monterey, CA, USA. ACM Press, New York, pp. 134–137.

Elmqvist, N., Fekete, J.-D., 2008. Semantic pointing for object picking in complex 3D environments. In: GI'08: Proceedings of Graphics Interface 2008, Windsor, Ontario, Canada. Canadian Information Processing Society, Toronto, Ont., Canada, pp. 243–250. ISBN:978-1-56881-423-0.

Elmqvist, N., Tsigas, P., 2006. View projection animation for occlusion reduction. In: Proceedings of the Working Conference on Advanced Visual Interfaces, Venezia, Italy. ACM Press, New York, pp. 471–475.

Elmqvist, N., Tsigas, P., 2008. A taxonomy of 3D occlusion management for visualization. IEEE Transactions on Visualization and Computer Graphics 14 (5), 1095–1109.

Flasar, J., Sochor, J., 2007. Manipulating Objects Behind Obstacles. In: Lecture Notes in Computer Science, vol. 4563. pp. 32–41.

Forsberg, A., Herndon, K., Zeleznik, R., 1996. Aperture based selection for immersive virtual environments. In: ACM UIST Symposium on User Interface Software and Technology, Seattle, Washington, USA. ACM Press, New York, pp. 95–96.

Grossman, T., Balakrishnan, R., 2004. Pointing at trivariate targets in 3D environments. In: ACM CHI Conference on Human Factors in Computing Systems. ACM, New York, NY, pp. 447–454.

Grossman, T., Balakrishnan, R., 2005. The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. In: ACM CHI Conference on Human Factors in Computing Systems, Portland, Oregon, USA. ACM Press, New York, pp. 281–290.

Grossman, T., Balakrishnan, R., 2006. The design and evaluation of selection techniques for 3D volumetric displays. In: ACM UIST Symposium on User Interface Software and Technology. ACM Press, New York, pp. 3–12.

Grossman, T., Wigdor, D., Balakrishnan, R., 2004. Multi finger gestural interaction with 3D volumetric displays. In: ACM UIST Symposium on User Interface Software and Technology, pp. 61–70.

Hinckley, K., Pausch, R., Goble, J.C., Kassell, N., 1994. A survey of design issues in spatial input. In: ACM UIST Symposium on User Interface Software and Technology. ACM, New York, NY, pp. 213–222.

Hwang, F., Langdon, P., Keates, S., Clarkson, P.J., 2003. The effect of multiple haptic distractors on the performance of motion-impaired users. Eurohaptics 14–25.

Keuning, H., 2003. Augmented force feedback to facilitate target acquisition in human–computer interaction. Ph.D. Thesis, Technical University of Eindhoven, Eindhoven.

Kim, S.-C., Kwon, D.-S., 2007. Haptic and sound grid for enhanced positioning in a 3-D virtual environment. In: Second International Workshop on Haptic Audio Interaction Design, pp. 98–109.

Lécuyer, A., Burkhardt, J.-M., Etienne, L., 2004. Feeling bumps and holes without a haptic interface: the perception of pseudo-haptic textures. In: ACM CHI Conference on Human Factors in Computing Systems. ACM, New York, NY, pp. 239–246.

Liang, J., Green, M., 1994. JDCAD: a highly interactive 3D modeling system. Computers and Graphics 18, 499–506.

Lindeman, R.W., 2003. Virtual contact: the continuum from purely visual to purely physical. In: Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society (HFES), pp. 2103–2107.

Livingston, M.A., Edward Swan II, J., Gabbard, J.L., Höllerer, T.H., Hix, D., Julier, S.J., Baillot, Y., Brown, D., 2003. Resolving multiple occluded layers in augmented reality. In: ISMAR'03: Proceedings of the Second IEEE/ACM International Symposium on Mixed and Augmented Reality. IEEE Computer Society, Washington, DC, USA, pp. 56–65. ISBN:0-7695-2006-5.

Looser, J., Billinghurst, M., Grasset, R., Cockburn, A., 2007. An evaluation of virtual lenses for object selection in augmented reality. In: Proceedings of the Fifth International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia, Perth, Australia. ACM, New York, pp.203–210.

McGuffin, M., Tancau, L., Balakrishnan, R., 2003. Using deformations for browsing volumetric data. IEEE Visualization, 401–408.

Meyer, D.E., Smith, J.E., Kornblum, S., Abrams, R.A., Wright, C.E., 1988. Optimality in human motor performance: ideal control of rapid aimed movements. Psychological Review 95, 340–370.

Mine, M., 1995. Virtual environment interaction techniques. UNC Chapel Hill Computer Science Technical Report TR95-018.

Oakley, I., McGee, M.R., Brewster, S., Gray, P., 2000. Putting the feel into look and feel. In: ACM CHI Conference on Human Factors in Computing Systems, pp. 415–422.

Oakley, I., Brewster, S., Gray, P., 2001. Solving multi-target haptic problems in menu interaction. In: ACM CHI Conference on Human Factors in Computing Systems (Extended Abstracts), pp. 357–358.

Oakley, I., Adams, A., Brewster, S., Gray, P., 2002. Guidelines for the design of haptic widgets. In: British HCI Conference, pp. 195–212.

Olwal, A., Feiner, S., 2003. The flexible pointer—an interaction technique for selection in augmented and virtual reality. In: ACM UIST Symposium on User Interface Software and Technology (Conference Supplement), Vancouver, BC, pp. 81–82.

Oviatt, S., 2002. Multimodal interfaces. In: The Human–Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications. Lawrence Erlbaum Associates, Inc., London, pp. 286–304.

Pierce, J., Forsberg, A., Conway, M., Hong, S., Zeleznik, R., 1997. Image plane interaction techniques in 3D immersive environments. In: ACM I3D Symposium on Interactive 3D Graphics, pp. 39–43.

Poupyrev, I., Billinghurst, M., Weghorst, S., Ichikawa, T., 1996. The go–go interaction technique: non-linear mapping for direct manipulation in VR. In: ACM UIST Symposium on User Interface Software and Technology, pp. 79–80.

Poupyrev, I., Weghorst, S., Billinghurst, M., Ichikawa, T., 1998. Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques. Proceedings of Eurographics, Computer Graphics Forum 17 (3), 41–52.

Ramos, G., Robertson, G., Czerwinski, M., Tan, D., Baudisch, P., Hinckley, K., Agrawala, M., 2006. Tumble! Splat! helping users access and manipulate occluded content in 2D drawings. In: Proceedings of the Working Conference on Advanced Visual Interfaces, Venezia, Italy. ACM Press, New York, pp. 428–435.

Rosenbaum, D., 1991. Human Motor Control. Academic Press, San Diego, CA.

Sallnäs, E.-L., Zhai, S., 2003. Collaboration meets Fitts' law: passing virtual objects with and without haptic force feedback. In: Proceedings of INTERACT, pp. 97–104.

Smyth, T.N., Kirkpatrick, A.E., 2006. A new approach to haptic augmentation of the GUI. In: Proceedings of the Eighth International Conference on Multimodal Interfaces, Banff, Alberta, Canada. ACM, New York, pp. 372–379

Steed, A., 2006. Towards a general model for selection in virtual environments. In: IEEE Symposium on 3D User Interfaces, pp. 103–110.

Steed, A., Parker, C., 2004. 3D selection strategies for head tracked and non-head tracked operation of spatially immersive displays. In: Eighth International Immersive Projection Technology Workshop, May 2004.

Unger, B.J., Nicolaidis, A., Thompson, A., Klatzky, R.L., Hollis, R.L., Berkelman, P.J., Lederman, S., 2002. Virtual peg-in-hole performance using a 6-DOF magnetic levitation haptic device: comparison with real forces and with visual guidance alone. In: Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. IEEE Computer Society, Silver Spring, MD, pp. 263–270.

Vanacken, L., Raymaekers, C., Coninx, K., 2006. Evaluating the influence of multimodal feedback on egocentric selection metaphors in virtual environments. In: First International Workshop on Haptic Audio Interaction Design, pp. 12–23.

Vanacken, L., Grossman, T., Coninx, K., 2007. Exploring the effects of environment density and target visibility on object selection in 3D virtual environments. In: IEEE Symposium on 3D User Interfaces, pp. 117–124.

Wall, S.A., Paynter, K., Shillito, A.M., Wright, M., Scali, S., 2002. The effect of haptic feedback and stereo graphics in a 3D target acquisition task. Eurohaptics, 23–28.

Wang, Y., MacKenzie, C., 2000. The role of contextual haptic and visual constraints on object manipulation in virtual environments. In: ACM CHI Conference on Human Factors in Computing Systems, pp. 532–539.

Ware, C., Franck, G., 1996. Evaluating stereo and motion cues for visualizing information nets in three dimensions. ACM Transactions on Graphics 15, 121–140.

Ware, C., Lowther, K., 1997. Selection using a one-eyed cursor in a fish tank VR environment. ACM Transactions on Computer-Human Interaction 4, 309–322.

Ware, C., Arthur, K., Booth, K.S., 1993. Fish tank virtual reality. In: ACM CHI Conference on Human Factors in Computing Systems, Amsterdam, The Netherlands. ACM Press, New York, pp. 37–42.

Wyss, H.P., Blach, R., Bues, M., 2006. iSith—intersection-based spatial interaction for two hands. In: IEEE Symposium on 3D User Interfaces, pp. 59–61.

Zhai, S., Buxton, W., Milgram, P., 1994. The "Silk Cursor": investigating transparency for 3D target acquisition. In: ACM CHI Conference on Human Factors in Computing Systems. ACM, New York, NY, pp. 459–464.