3D Selection Strategies for Head Tracked and Non-Head Tracked Operation of Spatially Immersive Displays

Anthony Steed, Chris Parker
Department of Computer Science, University College London
Gower St, London, WC1E 6BT, United Kingdom
A.Steed@cs.ucl.ac.uk, C.Parker@cs.ucl.ac.uk

ABSTRACT

We present an investigation of selection strategies designed specifically for the spatially immersive display (SID), CAVE™-like situation. We address two common contexts of use of SID displays: single user head tracked, and two person with one non-head tracked demonstrator and a visitor. This second context is, in our experience, a common use scenario of SIDs, but it, and other issues specific to SIDs, are largely ignored in the literature, which has concentrated on head-mounted displays. We introduce a novel interaction technique, Shadow Cone Selection, which was designed for the situation of non-head tracked interaction. Through analysis and experiment, we show that it might be preferred over standard cone and ray-based selection techniques for some applications.

KEYWORDS

3D Selection, 3D interaction, spatially immersive displays, virtual environments, usability

1. Introduction

Immersive virtual environment (IVE) systems including head-mounted displays (HMDs) and spatially-immersive displays (SIDs) such as CAVE™-like displays are finding increasing use in scientific and engineering domains [15]. In recent years, SIDs have become relatively prevalent and have eclipsed HMDs displays in many laboratories. However the associated domain of study of 3D interaction techniques is largely based around studies of HMD interaction and there are important differences between the HMD and SID situations.

Figure 1 shows a representation of typical SID in use by group of users. This is a typical context of use of a SID with two users inside the SID with others looking on. The first user is a visitor, who is head tracked and seeing the images in a first person view. They are accompanied by a demonstrator who controls the view through a hand-held device. The demonstrator only sees a distorted view of the world. There are also a pair of onlookers who also see a distorted view of the environment.

In these situations the requirements on interaction are very different from those presented in the literature. Most importantly, the demonstrator will have to point and select objects without having a first person view. This is somewhat similar to the situation of large-single plane displays, but an important difference is that the SIDs support and encourage near space interaction. In particular objects may appear to be within the SID space and within reaching distance.

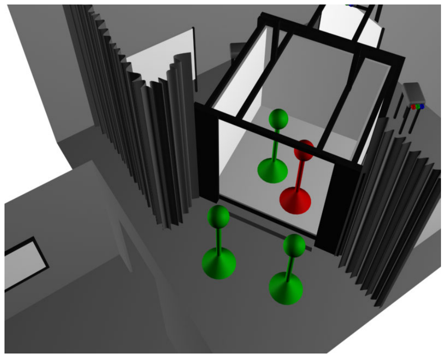

Figure 1. A representation of a group of users with a SID system. These include a client who is head-tracked, a demonstrator (in red) who accompanies them in the SID and a pair of onlookers

In this paper we look at selection metaphors for common situations of use of SIDs. We will focus on techniques that are easy to implement and are generally applicable across most common usage scenarios. We introduce a new selection technique, Shadow Cone, and in an experiment, compare it against two standard techniques, Ray Selection and Cone Selection [7]. In the following section, we introduce immersive interaction techniques and describe the difference between SID and HMD situations in more depth. We then give details of existing selection techniques and our new Shadow Cone technique. We then describe and discuss an experiment that compares these techniques in two common situations of use. Finally we will conclude by discussing some directions for further research in 3D interaction, and highlight some blind spots of current research.

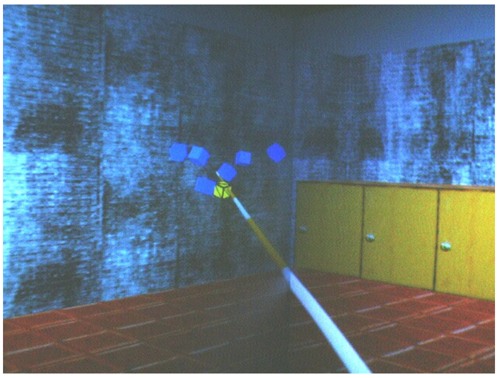

(a)

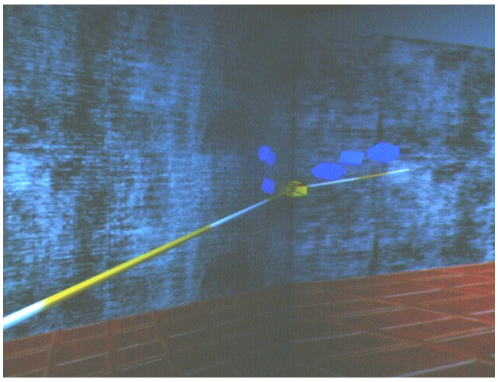

(b)

Figure 2. Two views from the SID of a user using Ray Selection to select the yellow object from a small cluster of distracting blue objects. (a) Shows a head tracked user's view of the scene. The ray appears straight. (b) Shows a non-head tracked user's view of the same situation where the view is generated for a head-tracked user who is standing 0.75 m to their right. Thus the projection is incorrect, and the ray appears bent across the corner between two SID walls. The non-head tracked user can still select the object.

2. Immersive Interaction

A complete review of current immersive virtual environment (IVE) interaction techniques is beyond the scope of this paper. Review material can be found in many papers in the area including [2,3,8,11]. Bowman et al. 2001 [3] provides a taxonomy of interaction tasks, including selection and manipulation. We will thus only describe a few critical issues in the choice and implementation of interaction techniques.

We will focus solely on selection in this paper. It is worth noting that selection is a necessary sub-task in manipulation but in practice, because selection can be done without the user's considering where the hand will have to move next, user strategies can be quite different. Furthermore, in the context of non-head tracked interaction, manipulation can be very difficult. In pilot experiments many subjects had severe problems undertaking manipulation tasks when not head tracked, so we decided not to pursue this at this time.

Selection within an IVE system is commonly effected by collision between the desired object and the user's virtual hand, or by intersection of the object with a ray projecting in the direction the hand is pointing. The former is more intuitive since it is analogous to the real world but it does require objects to be within reaching distance. The latter is more flexible but less natural since selection takes place at a distance.

The Go-Go interaction technique extends the virtual hand technique to support selection at a distance [10]. In [12] Poupyrev et al. note that Go-Go is a superset of the virtual hand and that whenever the virtual hand is used, Go-Go selection is a natural and flexible extension. Other selection techniques include image plane interaction technique [9] and aperture based selection [6].

2.1. Distinctions in User Experience for SID

Almost all of the techniques described in the previous section have been designed for and evaluated on HMD systems. The most important distinctive of SID systems is that they are multi-user systems but usually only one user can be tracked and the view is only correct for that user. The second user sees a distorted view of the view. Figure 2 shows a head tracked and non-head tracked view of the same scene. This suggests that there are two possibly conflicting goals in designing interaction techniques for a SID: to optimize for the tracked user, or to optimize with consideration of the comprehension of the group of users.

Another distinction is that the change from head-tracked to non-head tracked is very easy to effect. A group of users will often pass a hand controller around. Finally it is also worth noting that in our experience when demonstrating to large groups (four people or more) it is often helpful to disable head-tracking, at least temporarily, when explaining the basics of how the system works.

2.2. Evaluation of Interaction Metaphors

An overview of strategies for evaluation of virtual environments has been provided by Bowman, Gabbard, & Hix [4]. They discuss the difference between general usability engineering of applications and usability engineering of interaction techniques. Applications are used optimized to support a specific set of tasks, but interaction is usually optimized to support a generic class of tasks that support a wide range of applications. There is a tension between designing for a specific application and generalizing technique to other applications.

Bowman et al. [3] and Poupyrev et al. [11] have proposed test-beds for interaction tasks that contain a battery of standard interaction tasks. We have chosen not to take either of these test-beds, but rather to focus on a relatively

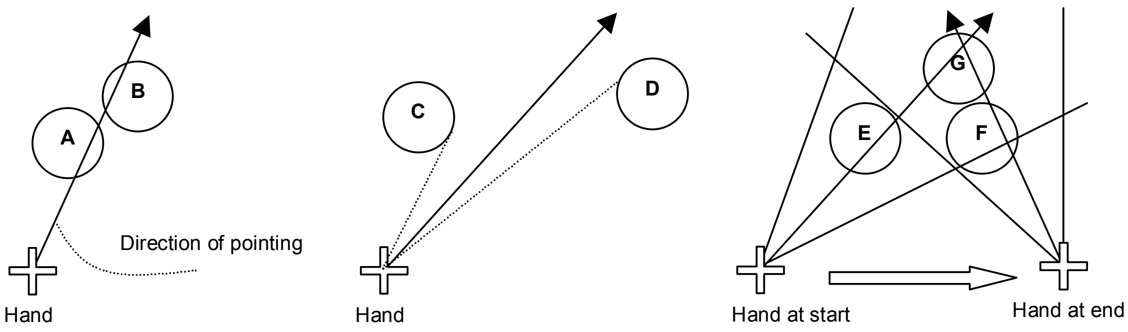

(a) (b) (c)

Figure 3. (a) Ray-based selection is effected by choosing first object to intersect the ray from the hand (Object A) (b) Cone Selection is effected by choosing the object that is relatively closest to the line (Object D). This is indicated by the angle to the dotted line to Object D subtending a smaller angle to the direction of pointing than the line to Object C (c) Shadow Cone Selection is effected by choosing the object that is within all the cones projected from the hand (Object G). In this case both Objects E and G are highlighted at the start, but E is dropped as the hand moves to its end position and direction.

small number of task types with a large number of repeated trials within each tasks. The main reason for this is that test-beds contain many tasks, and our main aim in this study is simply to highlight differences between the two cases of SID usage and three selection techniques.

Poupyrev et al. investigate the effects of distance, size and visual feedback on the performance of selection and manipulation tasks for virtual hand selection, ray selection and Go-Go hand selection [12]. They find that, as expected, performance decreases with greater distance and smaller size. They also find that visual feedback helps in some situations. The Go-Go hand was found to be a general extension of virtual hand. The Go-Go hand selection and ray selection performed roughly equally.

To our knowledge the only study of interaction specifically for SIDs is a study that compares SIDs & HMDs for selection or manipulation [16]. In that study it was found that although SIDs were superior for selection tasks, on manipulation tasks the systems were much closer. That study only compared the same techniques across the two conditions and does not offer any insight in to SID specific situations.

2.3. Interaction Guidelines

In order to aid design and evaluation, researchers have started to reify many of the informal insights and formal evaluations into guidelines. [5] is a recent collection of such guidelines that focuses on interaction. In the conclusions of this paper we will propose several new guidelines or clarifications to existing guidelines, based on observations in pilot experiments and main experiment.

3. Selection Techniques

We have chosen to focus solely on a few selection techniques. The main constraint will be that the techniques must be usable by non-head-tracked users.

SIDs support and encourage close by interaction for head tracked users because objects can appear within the space of the SID itself. However when this happens the two images of the object on the screen are drawn with negative parallax, and the closer to the center of projection that the object is, the greater the negative parallax. This parallax can get very large and it is not difficult to construct situations where left eye and right eye images appear on different walls of the SID. In such situations, the non head tracked user will be unable to fuse the image correctly unless they are extremely close to the head tracked user, and thus they will see a highly distorted view of the object. Objects that are close to and in front of or at any distance beyond the plane of the wall of the screen are only drawn with a small amount of parallax and are thus only slightly distorted for the non-head tracked user. Thus we will thus have to discount the virtual hand selection metaphor because it requires the user to intersect their hand to the geometry of the object to be selected. We also have to discount selection techniques such as Go-Go hand [10], image plane interaction [9] and aperture-based interaction [6] since they all require, or strongly prefer, an egocentric view.

In our situation manipulation at a large distance would be achievable using a technique such Go-Go hand [10] or Worlds in Miniature techniques [14]. Manipulation using ray-based techniques is also workable for objects near to or beyond the planes of the walls of the SID. However none of the techniques facilitates object placement in near space. In initial pilot trials even expert users had severe problems accomplishing placement tasks that head tracked users had no problems with. We note that this is a particular limitation of SIDs for non-head tracked users and it deserves further attention.

The three selection techniques we have chosen to study are Ray Selection, Cone Selection and a new technique Shadow Cone Selection.

3.1. Ray Selection

Ray Selection is one of the most frequently implemented selection techniques. The underlying technique of ray casting is commonly supported by most scene graph toolkits.

A ray is cast from the position of the hand in the direction of pointing of the hand (see Figure 3a). There are several possible interpretations of "direction of pointing", but most hand-held controllers have a natural and easily understood forwards direction. The controller that typically comes with an Intersense system is held in a single-handed tennis racquet style grip where it is the thumb that points forward. The closest object that intersects this ray is selected.

We have chosen to draw this ray (see Figure 2). We use an active selection mode. Once the user presses a button on the wand controller, the object intersecting the ray is highlighted. When the button is released the currently highlighted object is selected.

This technique is suitable for non-head tracked situations, because, as Poupyrev et al. note [12], it is effectively a 2D interaction technique. The user does not need to judge depth in order to select. Indeed despite the fact the non-head tracked user sees a distorted view and they might not be able to judge absolute angle horizontal and vertical angle to target, they can always rely on judgments of relative direction to target. Figure 2 showed a situation where the non-head tracked user is able to select an object from a small cluster.

3.2. Cone Selection

Ray Selection is difficult with small objects since the ray must intersect the object, and this can be especially difficult for the user if there is any presence of jitter in the tracking system. Cone Selection [7] alleviates this by using a metric that selects the object "closest" to a ray that is centered on the hand and points in the direction of the wand as described in the previous section (see Figure 3b). The measure of closeness is not simple distance to the ray, but the ratio of distance from the object to the closest point on the ray over the distance along the ray of the closest point. This can be roughly visualized by imaging holding the wrong (non-handle) end of an umbrella and then opening it. As it opens it sweeps a series expanding cones. The object selected is the one first intersected by the umbrella skin.

Again there are many implementation details. Cone Selection can be implemented by finding the object that subtends the smallest angle to the ray. There is usually a maximum angle that the object can subtend and usually this angle is indicated by drawing a transparent cone centered on the hand. To save calculating the distance to all objects in the scene, the geometry of this cone can be used to cull away objects that cannot be selected.

We have used an active selection mode. When the user presses the button, the closest object is highlighted and this is updated as the user moves their hand. When the user releases the button, the closest object is selected.

This technique is also suitable for non-head tracked users because it is more forgiving of failure to point accurately. Indeed it can be used in a searching style, when the cursor is moved around a group of objects without even intersecting the actual target object.

3.3. Shadow Cone Selection

Shadow Cone Selection is a novel selection technique. The previous two techniques select single objects. To select a group of objects, they can both be used in a mode where upon activation all objects that are intersected or are closest to the ray are selected. Shadow Cone is actually a group object selection technique, but it works in reverse. All the objects with a cone are selected, and as the user moves their hand only objects that are always within the cone are selected when the button is released (see Figure 3c).

To implement this, when the user presses the button, the potential selection set is initialized with all objects that are within a target angle of the ray from the hand. On each frame we remove any object is no longer within the target angle and we don't add new objects. When the user releases the button the objects remaining in the potential selection set are then selected. This can be visualized by imaging turning on a flashlight in a dim scene and selecting only the objects that are always in the beam of light before it is turned off.

This technique exploits the 3D nature of the display to a greater extent, because as Figure 3c shows, it allows selection to be specified by long gestures that can potentially express a complex spatial relation between the object and the hand.

4. Experiment

A between-subjects experiment was designed. The two independent variables were head tracking or non-head tracking, and Ray Selection, Cone Selection or Shadow Cone Selection. Thus there were six conditions. Thirty participants were recruited from students and staff at UCL. Each participant was paid £5 (approximately US$8) for taking part in the 40 minute experiment. Each participant undertook a set of six sub-tasks in the same order.

Before undertaking the six sub-tasks the user spent up to three minutes training to use the selection technique and mode of head tracking to which they had been randomly assigned. Training comprised of 5 selection trials of the type they were to undertake in the main experiment. The whole experiment took place in a simple virtual environment representing a room with a few pieces of furniture to aid orientation.

In the non-head tracked conditions, the viewpoint was fixed 1.7m above the center of the SID's floor. Locomotion was disabled throughout the experiment.

4.1. Sub-Tasks

Each sub task comprised 20 selection trials. Sub-tasks started easily and became progressively more difficult. In sub-task one, participants had to select the single large target object that appeared. In sub-task two, participants had to select the single small target object that appeared. Sub-tasks three and four comprised selecting the target object rather than a single distracting object. Task three used large objects whereas task four used small objects. Finally sub-tasks five and six involved selecting the target amongst a cluster of, respectively large and small distracting objects.

In every trial the user was instructed to stand in a new position. The position was indicated by a 0.5m by 0.5m red square on the ground. Note that the square was flat against the floor and thus, since it was coincident with a display wall, its image is correct for non-head tracked users as well as head tracked users. Thus there was no difficulty in standing on it. This part of the protocol was designed to simulate many situations of visual distortion for the user. The square appeared within a 2m by 2m area centered on the floor of the SID. This mean that the non-head tracked user was up to 1m from the center of projection.

Once the user stood on the square, the square disappeared and 1.0s later the target and any distracting objects appeared. Target objects were yellow. Distracting objects were blue. Active objects were highlighted by outlining in bold wire-frame. The objects appeared in a volume 5m deep by 7m wide by 1m high with the bottom edge 0.6m off the floor. The volume was in front of the back edge of the SID so the whole volume was visible to the user wherever they stood in the SID. Large objects were 0.32m along each edge. Small objects were 0.16m along each edge. In trials with pairs of objects, the objects were roughly 0.2m apart. In trials with clusters, the clusters comprised 7 objects where no object was more than 0.4m from the others.

Once the participant had selected the correct object the next trial started. If the participant took longer than 30 seconds, the trial was terminated and the next started. The time between a trial ending and the next red square being highlighted was 0.5s. There was no timeout for the user to stand on the square.

Between the training and the sub-tasks and between the sub-tasks the graphics were dimmed and a brief instruction given verbally by the experimenter.

4.2. Implementation

The SID used in these trials was a Trimension ReaCTor, with three back projected 2.8x2.2m walls, and a front projected 2.8x2.8m floor. This SID is typical of those used in the community. Visual output was produced by a SGI Onyx2 with eight 300MHz R12000 processors, 8GB ram and four InfiniteReality2 pipes. Stereo was achieved using shutter glasses. A full stereo image was generated for the user at 45Hz with the projectors running at 90Hz at a resolution of 1024x768.

The tracking system used was an Intersense IS-900. Users had the tracking unit attached to the top of the stereo glasses. Users held a wand controller with four buttons and an isotonic joystick. Only one button was used in the experiment.

The experimental system software was implemented in the VRJuggler framework version 1.1DR2 [1]. We used the OpenGL Performer library for rendering.

The software has a flexible system for configuring interaction techniques and implementing interaction experiments. Complete event logs were recorded so that each session can be replayed in its entirety.

5. Results

We present the results by the sub-tasks. To save space and aid readability, we use the following abbreviations in figures: Ray Selection (Ray), Cone Selection (Cone), Shadow Cone Selection (SCone), head tracked (TRK) and non-head tracked (N_TRK). When presenting statistical results we will give T-test values only in situations of relatively marginal significant and not for anything that is highly significant or is not significant.

We will be interested primarily in average selection times across the six conditions. We will also report whether or not the users failed to select the object within the timeout period, and how many times they selected the wrong object. To save space, we mention at the outset that with one exception for all sub-tasks and all selection tasks the head tracking case was strongly significantly faster than the non-head tracking case. The exception was Cone Selection for the Simple Small sub-task. Head tracked was still slightly faster than non-head tracked on average. We return to this in the Discussion section.

5.1. Simple Large

Average selection times for the Simple Large sub-task are shown in Figure 4. On no trial did any user select the wrong object, nor were there any timeouts.

There is no statistical difference between the three selection techniques for tracked users. For non-tracked users the only difference is between Cone Selection and Ray Selection.

5.2. Simple Small

Average selection times for the Simple Small sub-task are shown in Figure 5. For the head tracked conditions Shadow Cone Selection and Cone Selection were significantly faster than Ray Selection though there was no difference between

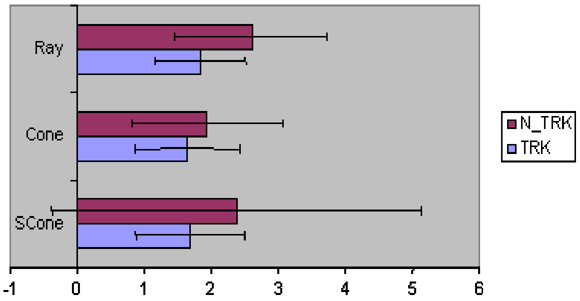

Figure 4. Average Selection time (s) for the Simple Large trials

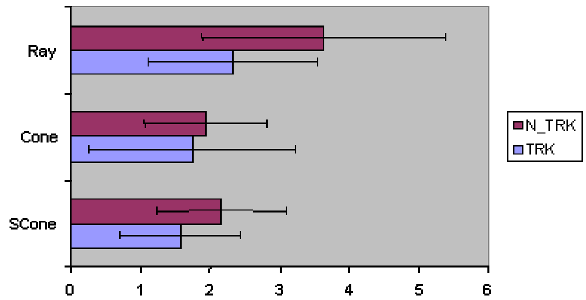

Figure 5. Average Selection time (s) for the Simple Small trials

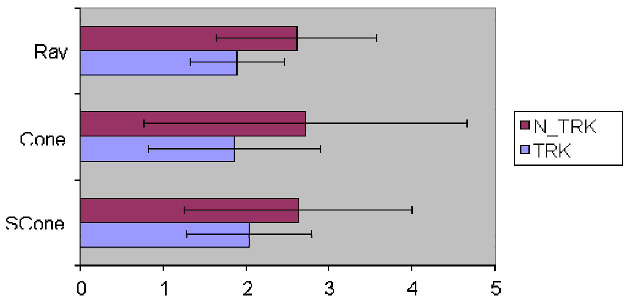

Figure 6. Average Selection time (s) for the Pair Large trials

Shadow Cone and Cone. The same occurs with the non-head tracked users.

There were no wrong selections. There were some timeouts. For the non-head tracked cases only, on average for each user there were for Ray Selection 0.4 timeouts, Cone Selection 0.2 and Shadow Cone Selection 0.4. This reflects the increased difficulty of the task.

5.3. Pair Large

In this sub-task there are no significant differences between any of the three selection techniques see Figure 6. However there are more wrong selections with the Cone Selection, with an average of 2.2 wrong selections per user on non-head tracked, 1.3 on head tracked. Shadow Cone Selection generated 0.3 wrong selections for non-head tracked and 0.2 for head tracked. There were no wrong selections with Ray Selection.

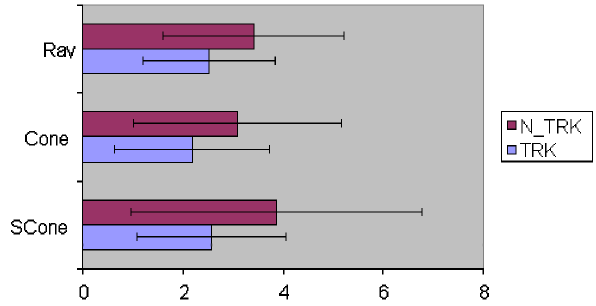

Figure 7. Average Selection time (s) for the Pair Small trials

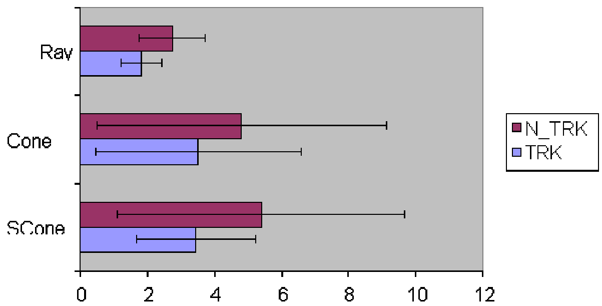

Figure 8. Average Selection time (s) for the Cluster Large trials

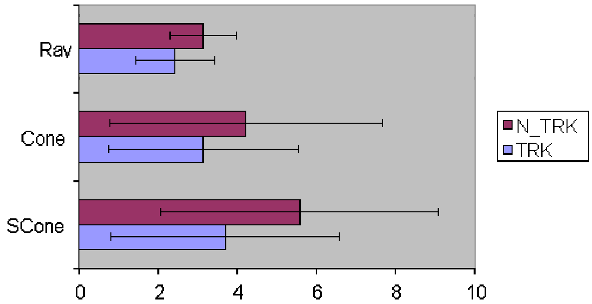

Figure 9. Average Selection time (s) for the Cluster Small trials

5.4. Pair Small

Cone Selection is just significantly faster that Shadow Cone Selection for both head tracked ($p<0.049$) and non-head tracked ($p<0.033$) (see Figure 7). There were no selection timeouts.

Cone Selection again shows more incorrect selections than the other two with an average of 2.7 wrong selections per user on non-head tracked, 0.6 on head tracked. Shadow Cone Selection generated 0.3 for non-head tracked and 0.2 for head tracked. Ray Selection generated 1.2 for head tracking, 0 for non-head tracking.

5.5. Cluster Large

Ray Selection is significantly faster than both Cone Selection and Shadow Cone Selection on both head tracked and non-head tracked (see Figure 8).

Cone Selection generated more timeouts with an average of 2.2 per user for non-head tracked, 1.3 for head tracked. Shadow Cone Selection generates 0.2 per user for non-head tracked only, and Ray Selection generated non. Cone Selection also generates more incorrect selections with 5.2 per user for non-head tracked, 20.3 for tracked. Shadow Cone Selection generated 7.8 per user on non-head tracked, 5.2 for tracked. Ray Selection generated 0.2 per user on non-head tracked, 0.5 for head tracked.

5.6. Cluster Small

Ray Selection is significantly faster than both Cone Selection and Shadow Cone Selection on both head tracked and non-head tracked (see Figure 9). On non-head tracked only, Cone Selection is faster than Shadow Cone Selection ($p<0.016$).

Cone Selection generates more timeouts with an average of 2.0 per user for non-head tracked, 2.0 for head tracked. Shadow Cone Selection generates 0.8 per user for non-head tracked only, and Ray Selection generated 0.2 per user for non-head tracked only. In this case Shadow Cone Selection generates most incorrect selections with 12.6 per user for non-head tracked, 4.5 for tracked. Cone Selection generated 4.4 per user on non-head tracked, 9.7 for tracked. Ray Selection generated 0.6 per user on non-head tracked, 1.2 for head tracked.

6. Discussion

6.1. Comparing Selection Techniques

Each of the selection techniques has a sub-task where it performs best. Cone Selection and Shadow Cone Selection are superior on the simple sub-tasks. The techniques are equal on tasks with pairs. Ray Selection is better in performance the sub-tasks involving clusters. On the clustered trials, we generally see many fewer timeouts and generally fewer incorrect selections with Shadow Cone selection. Overall then, there is no clear preference.

6.2. Comparing Head tracked and Non-head tracked

The differences between head tracking and non-head tracking are highly significant in all cases except one. Differences were Simple Large 0.58s, Simple Small 0.69s, Pair Large 0.72s, Pair Small 1.03s, Cluster Large 1.38s, Cluster Small 1.22s. That is, non-head tracked is between 34% and 47% slower. This is actually quite surprising, since as Figure 2 showed, the non-head tracked user actually sees quite a distorted view of the world. This result is encouraging for other types of large displays with groups of users that have a 1st person view and 1st person control metaphor but that don't use head tracking.

7. Conclusion

We have discussed 3D interaction metaphors for SID systems with three main objectives. First we have argued that interaction on SIDs is quite different than on HMDs and further design and evaluation of interaction metaphors needs to be undertaken. Second we have introduced a new interaction metaphor, Shadow Cone Selection. Finally we have examined in depth how this new technique compares with two common techniques, Ray Selection and Cone Selection for both head tracked and non-head tracked users.

Perhaps the most surprising result is that although non-head tracked users were slower than head tracked users in all conditions, the absolute difference was not high. Performance was between 34% and 47% slower on tasks taking 2-3 seconds. This suggests that more applications might take advantage of the two-user situation, with a pilot or demonstrator controlling the interaction and a mostly passive passenger having the correct view so they can benefit most from the immersive presentation.

Comparing across the three techniques reveals little to distinguish the three techniques on performance on the tasks chosen. Ray Selection is faster on some more complex tasks, whereas Cone Selection is faster on others. Shadow Cone Selection is slightly slower than Cone Selection but importantly it generates fewer timeouts, that is occasions when users fail to select any object, and also fewer incorrect selections.

The results of the experiment thus show that ray selection, cone selection and shadow cone selection all perform roughly equally. We have argued that interaction techniques for SID need to consider group usage scenarios. We noted in Section 3 a number of selection techniques that had to be rejected because they were simply not usable in the non-head tracked situation we can propose some new guidelines for 3D interaction:

- *Virtual hand should only be used for SIDs when the user will be head-tracked. We suggest that implementers avoid it for generic SID applications because of the variety of viewing scenarios.*
- *Ray, Cone and Shadow Cone selection are all suitable for non-head tracked SID users. Techniques that rely on an accurate hand-eye axis vector such as Go-Go hand, image plane selection or aperture selection are not.*
- *Close-by manipulation should only be attempted with head-tracked egocentric views. For the non-head tracked view, only manipulation at a distance should be attempted.*

It is worth noting that the Ray Selection and Cone Selection have limitations that can be quite severe. Cone selection is impractical in certain situations of selecting an object within a cluster that has a lot of occlusion of the target. Shadow Cone selection does not suffer so badly from this problem, and an expert should be able to select any object with few exceptions.

Ray Selection is poor when objects are very small. Cone Selection and Shadow Cone Selection do not suffer so much from problems of the size of the target. We have avoided both very small targets and highly complex situations after pilot trials where users became frustrated with these tasks. Thus we would propose a tentative guideline:

- *Shadow Cone Selection and Ray Selection are more suitable than Cone selection for highly complex selection tasks where the target is occluded.*

We have noted that Cone selection should be more resistant to jitter than Ray Selection because small jumps may not cause the highlighted object to be changed. The same is true of Shadow Cone Selection – small amounts of jitter are unlikely to drop objects. In this experiment we have not looked at jitter, and the Intersense IS-900 we are using generates very little perceptible jitter. However when developing the software platform we used a HMD equipped with a Polhemus ISOTRAK. This exhibited more jitter than the Intersense IS-900 and informally we confirmed that Cone or Shadow Cone Selection were slightly easier to use on this configuration. We can thus tentatively propose the following guideline:

- *Consider using Cone Selection or Shadow Cone Selection in preference to Ray Selection if there is any significant jitter in the tracking system.*

Although this can not be suggested as a guideline, we note Shadow Cone Selection has the property of consistently positively highlighting a superset of all potential objects that can be selected by the current selection action. Any sudden movement or jitter when finishing the selection action can only result in a potentially null subset of the desired set being selected. This is in contrast to Ray Selection and Cone Selection where sudden movement or jitter at the end of the selection action can lead to a totally different object or set of objects being selected.

Finally we would note that Shadow Cone Selection appears to be novel and it sits in a potential blind spot of current interaction taxonomies. This suggests that there is still a lot of interesting work to be done in exploring the space of interaction metaphors for the expanding range of immersive display devices.

ACKNOWLEDGMENTS

Thanks to David Swapp for help on UCL's ReaCTor system. All the code, data and documentation necessary to run the experiments are available on the web at http://www.cs.ucl.ac.uk/research/vr/Projects/Interact/

REFERENCES

1. Bierbaum, A., Just, C., Hartling, P., Meinert, K., Baker, A. &Cruz-Neira, C. VR Juggler: A Virtual Platform for Virtual Reality Application Development, *Proc. of IEEE VR2001*
2. Bowman, D. & Hodges, L. Formalizing the Design, Evaluation, and Application of Interaction Techniques for Immersive Virtual Environments. *Journal of Visual Languages and Computing*, 1999, pp. 37-53.
3. Bowman, D., Johnson, D. & Hodges, L. Testbed Evaluation of Virtual Environment Interaction Techniques. *Presence: Teleoperators and Virtual Environments*, 10(1), 2001, pp. 75-95.
4. Bowman, D., Gabbard, J. & Hix, D A Survey of Usability Evaluation in Virtual Environments: Classification and Comparison of Methods. *Presence: Teleoperators and Virtual Environments*, 11(4), 2002, pp. 404-424.
5. Bowman, D. Principles for the Design of Performance-oriented Interaction Techniques. *In Handbook of Virtual Environments*, Stanney, K. (ed.), Lawrence Erlbaum Associates, 2002, pp. 277-300
6. Forsberg, A., Herndon, K., Zeleznik, R. Aperture Based Selection for Immersive Virtual Environments, *Proc. UIST'1996*, pp. 95-96.
7. Liang, J., and Green, M. JDCAD: A highly interactive 3D modeling system. *Computers & Graphics*, 18(4), 1994, 499-506.
8. Mine, M., Brooks Jr., F. P. & Sequin, C. Moving Objects in Space: Exploiting Proprioception in Virtual-Environment Interaction. *Proc. of SIGGRAPH 97*.
9. Pierce, J., Forsberg, A., Conway, M., Hong, S., Zeleznik, R., and Mine, M. Image Plane Interaction Techniques in 3D Immersive Environments. Proc. 1997 Symp. on Interactive 3D Graphics, pp 39-44.
10. Poupyrev, I., Billinghurst, M., Weghorst, S. & Ichikawa, T. Go-Go interaction technique: Non-linear mapping for direct manipulation in VR. *Proc. of UIST'96*. ACM. pp. 79-80.
11. Poupyrev, I., Weghorst, S., Billinghurst, M. & Ichikawa, T. A Framework and Testbed for Studying Manipulation Techniques for Immersive VR. *Proc. ACM VRST'97*.
12. Poupyrev, I., Weghorst, S., Billinghurst, M. & Ichikawa, T. Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques. *Computer Graphics Forum 17(3)*, 1998, pp. 41-52.
13. Stanney, K.M., Mourant, R. & Kennedy, R.S. Human factors issues in virtual environments: A review of the literature. *Presence: Teleoperators and Virtual Environments*, 7(4), 1998, 327-351.
14. Stoakley, R., Conway, M., Pausch, R. Virtual Reality on a WIM: Interactive Worlds in Miniature. *Proc. SIGCHI '95*, pp. 265-272.
15. Stone, R.J. Applications of Virtual Environments: An Overview, *In Handbook of Virtual Environments: Design, Implementation and Applications*, Stanney, K.M (editor), Lawrence Erlbaum Associates, 2002, pp. 827-856.
16. Steed, A. Parker, C. Evaluating Effectiveness of Interaction Techniques across Immersive Virtual Environment Systems, *submitted for publication.*