Remote Object Translation Methods for Immersive Virtual Environments

J.D. Mulder

Remote Object Translation Methods for Immersive Virtual Environments

Jurriaan D. Mulder

CWI,

P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

ABSTRACT

In this paper, seven methods are described to perform remote object translations with a six degree-of-freedom input device in an immersive virtual environment. By manipulating objects remotely, a number of disadvantages of the real-world 'direct grab and drag' metaphor can be avoided. The different methods are evaluated with a pilot user experiment. From the results of the experiment, some initial recommendations are formulated on the use of the methods for different manipulation tasks.

1. INTRODUCTION

Virtual reality systems are designed to provide the effect of *immersion* in an *interactive* three-dimensional computer-generated environment [2]. Virtual environments in potential can provide more intuitive methods for viewing and user interaction. A wide variety of 3D graphics applications could benefit from these methods. However, whereas the viewing aspect of virtual environments is well understood, user input is still a current research issue.

User input in virtual environments is usually performed with multiple degree-of-freedom input devices. A survey of various research results in this area of *spatial input* was composed by Hinckley et al. [5]. More recent results on methods for 3D object manipulation with multiple degree-of-freedom input devices are for instance given by Zhai et al. [12] who report on human performance in six degree-of-freedom input, Cutler et al. [4] who report on two handed interaction, and Bowman et al. [1] who evaluated different techniques for grabbing and manipulating remote objects. Most of these methods however, are designed to incorporate multiple types of transformations, primarily combined rotation and translation.

In this paper, we describe seven different methods for object translations only, and present the results of a pilot user experiment conducted to evaluate the methods on the aspects of speed, accuracy, and user friendliness. Although we have focussed on the use of these methods in a CAVE [3] equipped with the hand-held six degree-of-freedom input device (the *wand*), the results are applicable to other immersive virtual environments with different input and display devices as well.

2. WHY REMOTE?

As the CAVE is an immersive environment, an obvious manipulation scheme is the real-world metaphor: direct grab and drag. Here, the user moves his hand with the wand to the position of the object and 'grabs' the object by pushing a wand button. The user then drags the object to a new position as if he holds the object in his hand. This approach is both simple and intuitive. However, it does have four major disadvantages.

First of all, when the physical hand is moved onto the object, the hand and arm of the user are brought into the area onto which the user's view is focused. In the CAVE, this causes a disturbance of the correct view of the scene: objects that are (virtually) located in between the user's eyes and hand are no longer viewed correctly. As a result, the user has great difficulty determining whether his hand is at the same position as the object to

be picked. Once the object is picked, exact positioning of the object becomes difficult for the same reason. A similar (albeit reverse) violation of the occlusion cue was reported by Schmandt [8] as a major factor for incorrect depth judgements.

The second disadvantage is that the device that tracks the wand's position must be very accurate. Small discrepancies between the physical and virtual position of the wand become very visible and disturbing as the user is directly watching the wand and the object and expects an absolute correspondence between the positions of the two. Such distortions become less of a problem when relative positioning can be used [10], e.g., when the wand and object are not at the same location and there is an indirect connection between the object and the user's hand.

The third disadvantage is that the user can only manipulate those objects that are within reach. The user's arms are of limited length, and the CAVE has only a limited space in which the user can maneuver, whereas the virtual world in principle can be infinitely large. To manipulate objects that are not within reach, or to position objects to locations that are out of reach, the user has to perform a navigation action to move himself through the virtual world. Bowman et al. [1] state that navigation is desirable, as it allows different visual perspectives, but it should not be required.

Finally, as also reported by Bowman et al. [1], a direct 'real-world' grab and drag is cumbersome for the manipulation of large objects. Frequent navigation is required as such objects occlude the user's view if they are within reach of the user's arm, thereby preventing the user from obtaining a proper perception of the current position of the object in the scene.

All of these disadvantages to a great extent are directly related to the real-world manipulation metaphor and can be avoided by the use of techniques for *remote* object translations, where the objects are not required to be within reach of the user.

3. TRANSLATION METHODS

Before an object can be translated remotely it has to be picked. Bowman et al. [1] report that ray casting is a good method for distant object picking. With ray casting, the wand emits a visible ray. The user can point the ray such that it passes through an object. Pressing the left wand button picks the object hit by the ray (if the ray passes through multiple objects, the object closest to the user is picked). The object can now be manipulated. Releasing the wand button releases the object which then stays at its current position, i.e., the button serves as a clutch.

We have evaluated seven different methods for remote object translations. The methods can be classified into different categories according to two important characteristics: Which degrees of freedom of the wand are used to perform the object translations, and the mapping scheme from the wand motions to the actual object translations. For the first characteristic, a distinction is made between methods that are related to (changes in) the wand *position, orientation*, or a *combination* of both. For the second characteristic a distinction is made between methods that use object *position* control, object *velocity* control, or a *combination*. Table 1 depicts the classification of each method according to the forementioned characteristics.

		Wand Parameters		
		position	position + orientation	orientation
Motion Control	position	Slave	Stick Crosshair	
	position + velocity	Slave & Fly	Stick & Throttle	
	velocity	Fly		Throttle

Table 1: Classification of the manipulation methods according to their characteristics on degree of freedom dependency (horizontal) and object motion control (vertical).

A short description of each of the methods that were implemented is now given. Each method is also illustrated in a figure. This figure depicts how an example translation can be performed with the particular method. A sphere is to be translated from a start position to a goal position. The example translation is shown in figure 1 at the instant the sphere is picked for translation.
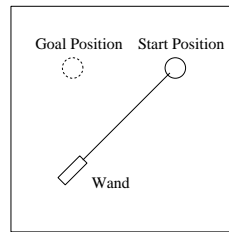

Figure 1: Example translation task at the instant the sphere is picked for translation.

3.1 Position Control

Three different methods were investigated that use position control as the mapping function from the wand's movements to the objects translations. The first method is solely based on the position (changes) of the wand, whereas the other two methods incorporate both the wand's position and orientation parameters.

Slave This method is based on a direct, one-to-one position mapping between the wand and the object: each translational movement of the wand is also applied to the object. This method is referred to as the 'Slave' method as the relation between the wand and the object can be seen as a master-slave relation: the object follows the exact (translational) movements of the wand. A disadvantage of this method is that the amount of translation is bounded by the reach of the user's arm. For larger translations, the concept of *ratchetting* has to be used, i.e., a series of grab-translate-release actions. The Slave method is illustrated in figure 2.
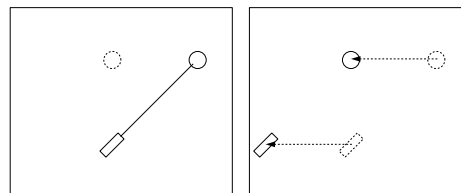

Figure 2: The Slave method. The sphere simply follows the wand movement.

Stick In this method the object is connected to the ray emitted by the wand upon the grab action. Then, each translational and rotational movement of the wand is also applied to the object, i.e., the wand is used as a pointer stick to position the object. Therefore this method is referred to as the 'Stick' method. Again, ratchetting has to be used for large translations, although in this case only for large object translations towards or away from the user. The Stick method is illustrated in figure 3.

3D Crosshair This method is based on a crosshair cursor method that has been employed successfully in 3D graphics applications for desk-top computers where 3D positioning is performed with a conventional 2D mouse [7]. Here, a 3D crosshair cursor is projected on the 2D screen. The user can translate the crosshair by dragging the 2D mouse pointer along one of the projected axes of the 3D crosshair. The same scheme is employed in the CAVE version of this method: when the object is picked, a 3D crosshair
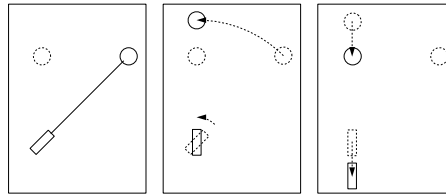
Figure 3: The Stick method. The sphere moves as if attached on a stick.

cursor appears at the intersection point of the ray and the object. The user can now translate the object by dragging the ray of the wand along one of the crosshair axes. In this method, no ratchetting is needed to perform large object translations. However, disadvantages are that only translations along the crosshair's axes can be performed, and no translations can be performed along axes of the crosshair that are parallel to the wand's ray. The Crosshair method is illustrated in figure 4.
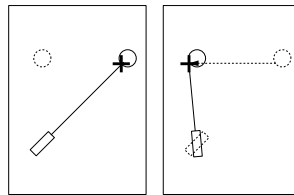

Figure 4: The Crosshair method. The crosshair (and the sphere) is moved by pointing the wand ray along one of the axes of the crosshair.

3.2 Velocity Control
Two methods were investigated that are based on velocity control. The first method is based solely upon the wand's position, while the other method is solely based on the wand's orientation.

Fly This method maps the change in position of the wand to a direction and velocity of a motion that is applied to the object. The position of the wand at the instant the object is picked is considered the rest position. When the wand is moved to a different position, the difference vector from this new position to the rest position indicates the direction of translation. The velocity of the object is derived from the length of the difference vector: the longer the vector, the faster the object moves. Below a certain threshold length no motion is applied to the object. To halt the object, the wand is brought back to the rest position. This method is referred to as the 'Fly' method as the user is in control of what can be regarded as the object's flight. The Fly method is illustrated in figure 5.
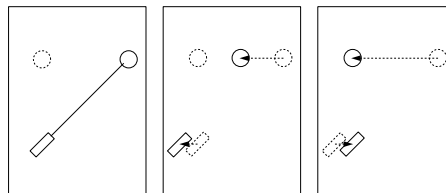

Figure 5: The Fly method. A motion is applied to the sphere by moving the wand away from its rest position. To stop the sphere the wand is brought back to the rest position.

Throttle This method is based on the orientation of the wand. The user indicates the direction of the motion by pointing the wand (ray). To control the velocity of the object, a motorcycle throttle grip metaphor is used: by rotating the wand about the direction of the motion (i.e., the *roll* of the wand) the velocity of the motion is indicated. A counter clockwise rotation implies a forward motion along the indicated direction, a clockwise rotation a backward motion. In both directions, a threshold value on the amount of rotation applies below which no translations are performed. To halt the object, the wand is rotated back to the rest orientation. This method is referred to as the 'Throttle' method. It is illustrated in figure 6.
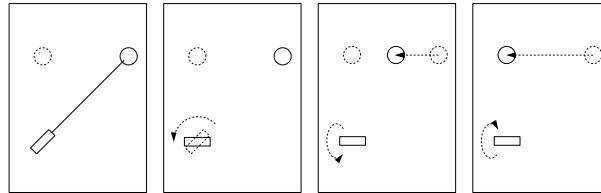

Figure 6: The Throttle method. The direction of movement is indicated by the wand. A velocity is implied on the sphere by rotating the wand about the direction of movement. To stop the sphere the wand is rotated back to its rest orientation.

3.3 Hybrid Control

Previous research results have shown that object positioning with velocity control is often less precise than with position control (e.g., [11, 6]). Therefore, two hybrid methods were created that use both velocity and position control. The intention of these methods is to combine the benefits of velocity control (facile large object translations) with the possibly more precise object positioning obtained with position control.

Slave & Fly In the Fly method, there is a spherical volume around the wand's rest position where no motions are applied to the object. The radius of this sphere equals the threshold value used for the derivation of the velocity from the length of the difference vector between the current wand position and the rest position. In this hybrid method, wand motions inside the sphere are directly coupled to the object position, i.e., the Slave method is used inside the sphere. The Slave & Fly method is illustrated in figure 7.
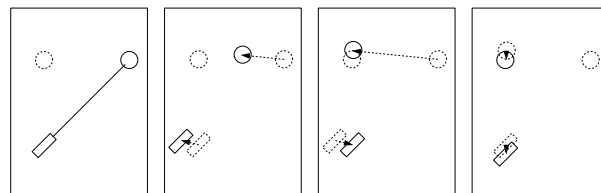

Figure 7: The Slave & Fly method. If the sphere is not translated to the exact desired location with the Fly method, small corrections can be made according to the Slave method.

Stick & Throttle The second hybrid method is a mix of the Stick and Throttle methods. Actually, it is the same as the Stick method, only now the throttle metaphor has been added for translations along the stick, i.e., the object can be moved towards and away from the user by rotating the wand about the roll axis. In [1] such control over the distance from the object to the user is referred to as a *fishing reel* metaphor. The Stick & Throttle method is illustrated in figure 8.
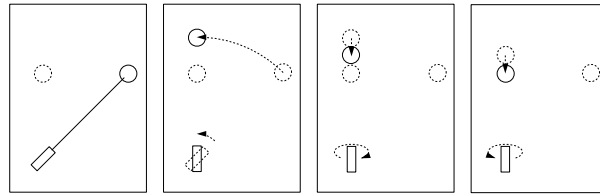
Figure 8: The Stick & Throttle method. The sphere can be translated towards and away from the wand with the throttle metaphor.

4. EVALUATION

4.1 Experiments

To test the manipulation methods on their usability, a small pilot experiment was conducted. Four test subjects were asked to complete a number of object translation tasks. Each of these subjects was familiar with the CAVE, but none had previously used any of the object manipulation techniques. While standing in a square virtual work space of equal size as the physical CAVE, a sphere was to be translated from a start position onto a goal position inside a bounding box. The sphere was constrained to stay inside the workspace. To provide additional feedback on the 3D position of the objects, orthogonal shadow projections were depicted on each of the side planes of the working space (i.e. the walls of the CAVE).

First, the different methods were explained to the test subjects and they were given a practice session to get familiar with the methods. Then, for each method a session of eight diffrent translation tasks had to be completed. The subjects were asked to bring the sphere from its start position to the goal position.

The experiments were conducted to evaluate the manipulation methods on three criteria: accuracy, speed, and user friendliness. To be able to determine the accuracy and speed of a method, all movements of the sphere invoked by the test subjects were logged to file. In the analysis, the first trial position performed with each method was discarded. User friendliness is a subjective merit. It was evaluated by having the test subjects fill out a small questionnaire about the different methods, and by making observations of the subjects while performing the manipulation tasks.

4.2 Results

The number of subjects and the sample size is small. Therefore, a full statistical analysis is not appropriate. The tests are primarily intended as a pilot study to obtain indications on the usability of the techniques and to evaluate the current implementations and parameter settings of the techniques. The experiences obtained from these experiments can be used for the design and conduction of a more formal and extensive usability analysis. Nevertheless, several interesting observations can be made.

Figure 9 shows box and whisker plots [9] of the task completion times per method and per subject. The figure shows great variety among the different methods. From this figure, a ranking can be obtained where the Stick method is ranked best, followed by the Stick & Throttle method, the Slave, Slave & Fly, and Fly methods, and finally the Throttle and Crosshair methods. The Throttle and Crosshair methods perform really bad, and reveal some unacceptable outside and far outside measurements.

Figure 9 also shows some variety among the subjects. If a ranking was to be made, subject 2 would score best, followed by subjects 3 and 1, and finally subject 4. Each of the subject plots reveal some outside or far outside scores. All of these scores were obtained with either the Throttle method or the Crosshair method.

To investigate the methods more thoroughly on the two aspects of speed and accuracy, the completion time was split in two components. The first is the time it took for the subjects to bring the sphere into the vicinity of the box, i.e. the point in time where the center of the sphere was located inside the box for the first time. This component can be regarded as a measurement for manipulation speed, i.e. how rapidly can the sphere be brought into the vicinity of the box with only limited accuracy. The second component is the time it took for the subjects to bring the sphere from that 'vicinity position' to the final position. This can be regarded as a
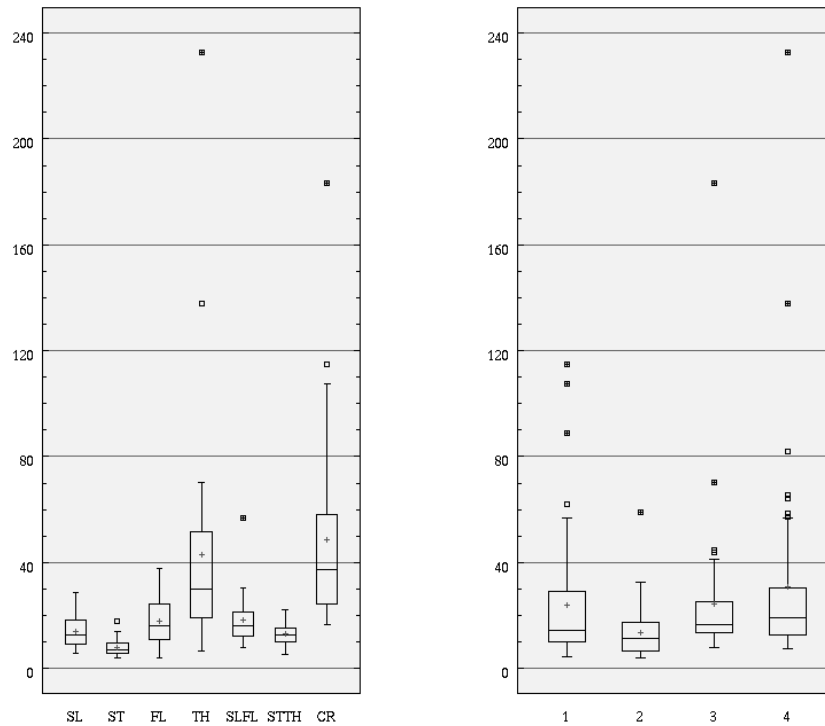
Figure 9: Box plots of the overall completion time in seconds of all trials. Left per method (SL = Slave, ST = Stick, FL = Fly, TH = Throttle, SLFL = Slave & Fly, STTH = Stick & Throttle, CR = Crosshair), right per subject.

measurement for accuracy, i.e. how easily can the final adjustments be made to accurately position the sphere inside the box. Besides looking at the final completion time of the trials as a measure for accuracy, the final distance of the sphere to the ideal end position was also considered.

4.3 Discussion

From the pilot user experiments it could be concluded that the Stick method performs best overall. The subjects perceived it as easy to use, intuitive, and not fatiguing. Furthermore, the obtained measurements indicated that the Stick method performs well for both fast and precise object placement. However, the object translation tasks as they were performed during these experiments did not require the sphere to be translated over large distances towards or away from the subjects. If a manipulation method is required to perform such large scale object translations, the Stick & Throttle method would be a better choice as far as fast object translations are concerned. However, precise object placement performance with the Stick & Throttle method might degrade as the object is located far away from the subject. In such a case, even small wand movements will result in large object translations. The Slave & Fly method might be a better choice in such a case. Further and more extensive experiments are needed to be conclusive about this. One thing was clear though: the Throttle and Crosshair methods performed by far the worst, and are not likely to be useful to perform object translations in a virtual environment.

5. CONCLUSION

User input is an underdeveloped aspect of the use of virtual environments. In particular, spatial input techniques for 3D object manipulations with the use of multiple degree-of-freedom devices is still a current research issue. We have evaluated seven different methods for remote object translations in immersive environments. The

results of the pilot user experiments showed a wide variety among the different methods, and some initial recommendations could be formulated on the use of the methods for different tasks. We plan to perform more extensive user experiments to be conclusive about these recommendations.

References

1. D.A. Bowman and L.F. Hodges. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In S.N. Spencer, editor, *1997 Symposium on Interactive 3D Graphics*, pages 35–38, 1997.

2. S. Bryson. Approaches to the successful design and implementation of VR applications. In R.A. Earnshaw, J.A. Vince, and H. Jones, editors, *Virtual Reality Applications*, pages 3–15. Academic Press, 1995.

3. C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 135–142, 1993.

4. L.D. Cutler, B. Fröhlich, and P. Hanrahan. Two-handed direct manipulation on the responsive workbench. In S.N. Spencer, editor, *Proceedings of the 1997 Symposium on Interactive 3D Graphics*, pages 107–114, 1997.

5. K. Hinckley, R. Pausch, J.C. Goble, and N.F. Kassell. A survey of design issues in spatial input. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '94)*, pages 213–222, 1994.

6. W.S. Kim, F. Tendick, S.R. Ellis, and L.W. Stark. A comparison of position and rate control for tele-manipulators with consideration of manipulator dynamics. *IEEE Journal of Robotics and Automation*, RA-3(5):426–436, October 1987.

7. J.D. Mulder and J.J. van Wijk. 3D computational steering with parametrized geometric objects. In G.M. Nielson and D. Silver, editors, *Visualization '95 (Proceedings of the 1995 Visualization Conference)*, pages 304–311, 1995.

8. C. Schmandt. Spatial input/display correspondence in a stereoscopic computer graphic work station. *Computer Graphics*, 17(3):253–261, 1983.

9. J.W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

10. C. Ware. Using hand position for virtual object placement. *Visual Computer*, 6(5):245–253, 1990.

11. S. Zhai and P. Milgram. Human performance evaluation of manipulation schemes in virtual environments. In *Proceedings of the 1993 IEEE Virtual Reality Annual International Symposium*, pages 155–161, 1993.

12. S. Zhai, P. Milgram, and A. Rastogi. Anisotropic human performance in six degree-of-freedom tracking: An evaluation of three-dimensional display and control devices. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*, 27(4):518–528, 1997.