

Voodoo Dolls: Seamless Interaction at Multiple Scales in Virtual Environments

Jeffrey S. Pierce, Brian C. Stearns, Randy Pausch
Carnegie Mellon University
{jpierce, bstearns, pausch}@cs.cmu.edu

Abstract

The Voodoo Dolls technique is a two-handed interaction technique for manipulating objects at a distance in immersive virtual environments. This technique addresses some limitations of existing techniques: they do not provide a lightweight method of interacting with objects of widely varying sizes, and many limit the objects that can be selected and the manipulations possible after making a selection. With the Voodoo Dolls technique, the user dynamically creates *dolls*: transient, hand held copies of objects whose effects on the objects they represent are determined by the hand holding them. For simplicity, we assume a right-handed user in the following discussion. When a user holds a doll in his right hand and moves it relative to a doll in his left hand, the object represented by the doll in his right hand moves to the same position and orientation relative to the object represented by the doll in his left hand. The system scales the dolls so that the doll in the left hand is half a meter along its longest dimension and the other dolls maintain the same relative size; this allows the user to work seamlessly at multiple scales. The Voodoo Dolls technique also allows both visible and occluded objects to be selected, and provides a stationary frame of reference for working relative to moving objects.

CR Categories and Subject Descriptors: I.3.6 [Computer Graphics] Methodology and Techniques - Interaction Techniques, I.3.7 [Computer Graphics] Three-Dimensional Graphics and Realism - Virtual Reality.

Additional keywords: virtual worlds, head-mounted display, two handed interaction, bimanual frame of reference

1 INTRODUCTION

Virtual reality allows users to interact with objects in a virtual world that are beyond their physical reach. The problem is to create an interaction technique that allows interaction with any object in the world. Existing solutions do not provide a lightweight method of interacting with objects of widely varying sizes. For example, the Worlds In Miniature technique [19] provides the user with a miniature replica of the world that he can manipulate to affect the larger world, but shrinking the entire world to a small, hand-held object can make it impossible for the user to select small objects and place them precisely. Many of these solutions also limit the




Figure 1: Manipulating a pin and toy soldier with dolls held in the dominant and non-dominant hands. In all figures we added the user's hands to the image for increased clarity.

objects that can be selected or the manipulation possible after making a selection. For example, the Go-Go interaction technique [16] scales the motion of the user's hands to extend his reach, but still limits the objects a user can select and where he can put them by how far he can extend his arms. Because of these limitations, interaction with objects at a distance is still an unsolved problem.

We propose the Voodoo Dolls interaction technique to address these limitations. This technique is a two-handed interaction technique for immersive virtual environments. Our technique:

1. Allows the user to work seamlessly at multiple scales.
2. Allows both visible and occluded objects to be manipulated.
3. Takes advantage of research showing that a user's dominant hand works in the reference frame defined by his non-dominant hand [7,9].
4. Simplifies working relative to moving objects.

2 VOODOO DOLLS TECHNIQUE

Dolls are hand held, transient objects: users create dolls to represent specific objects, and the dolls vanish when they are no longer needed. Visually, dolls are small copies of the objects they represent. Functionally, dolls are used in pairs, one in each hand, and serve two different yet complementary purposes depending on the hand holding them.

Guiard [7] showed that the dominant hand works in the reference frame defined by the non-dominant hand. Consider signing a baseball: the baseball is much easier to sign when held in the non-dominant hand than when sitting in a stand on a table. We take advantage of this by making the effects and properties of the dolls depend on whether the user holds them with his dominant or non-dominant hand. To simplify the following discussion, we assume that the user is right-handed (the right hand is the dominant hand). The doll held in the right hand determines the position and orientation of the object it represents, while the doll held in the left hand sets the reference frame. In other words, moving the doll in the right hand relative to the doll in the left hand moves the object represented by the doll in the right hand to the same position and orientation relative to the object represented by the doll in the left hand. This allows the user to quickly and easily position objects relative to each other in the virtual world. Consider a user who wants to put a lamp on a table. The user creates dolls for both objects, and holds the table doll in his left hand and the lamp doll in his right hand. Now, placing the lamp doll on top of the table doll causes the lamp to move on top of the table.

2.1 Creating Dolls

The user creates a doll using an image plane technique [15]: he positions his hand so that the crosshair attached to that hand appears over an object in his line of sight and pinches together the index finger and thumb of that hand (see Figure 2). This selection is analogous to placing the mouse cursor over an object on a desktop display. Selecting an object creates a doll that is scaled to a convenient working size (we use half a meter along its longest dimension) and placed in the user's hand.



Figure 2: Creating A doll

Users can pass dolls from hand to hand. For example, to pass a doll from the right to the left hand, the user moves the cursor attached to his left hand to the doll and pinches the thumb and forefinger of that hand together; releasing the right hand's pinch now transfers control of the doll to the left hand and changes the doll's mode. When the user releases all pinches holding a doll, the doll vanishes. Any accumulated manipulations are retained when dolls are passed from hand to hand. We provide a limited form of undo using Mine's "over the shoulder" deletion technique [13]: the

user can reset an object's position and orientation by holding the doll that represents it over his shoulder and releasing it.

The following example illustrates the Voodoo Dolls technique. Imagine a theatrical director using virtual reality techniques to experiment with different layouts of set furniture for a stage play. He is sitting thirty rows back in a virtual theater, and wants to move a telephone onstage from a desk to a coffee table. If the director is right handed, he creates a doll for the desk by "grabbing it" with his left hand. To help provide context (see section 2.3), the system creates dolls for the objects on the desk, which appear on the desk's doll in the director's left hand. With his right hand, the director reaches into the context and grabs the doll for the telephone. He releases his left hand's grip, destroying the desk doll and its context, then creates a doll for the coffee table with his left hand. He places the telephone doll on the coffee table doll, and lets go of both. During this process, whenever both his hands are holding dolls the onstage object represented by the doll in his right hand is positioned relative to the onstage object represented by the doll in his left hand.

2.2 Scaling

In addition to changing the effects of dolls based on the hand they are in, we also change their properties. We determine the size of dolls based on the hand holding them. Our goal is to provide a lightweight mechanism for the user to change the scale he is working at: a user designing a virtual town should be able to move a house to the other side of town, then put a dog next to a doghouse, and finally place a bee on a flower all without explicitly changing scale. We accomplish this in the Voodoo Dolls technique by scaling the doll in the user's left hand to be half a meter along its largest dimension, and then resizing the other dolls to maintain their relative size. Thus the dolls are a comfortable working size regardless of the size of the actual objects they represent. Note that this does require choosing an appropriate reference object to be most effective: choosing the house that the flower is next to as the reference object instead of the flower itself will make the doll for the bee much smaller and harder to manipulate.

2.3 Providing Context

When the user holds a doll in his left hand, the system can create dolls for nearby objects to provide some context for the interaction. These new miniatures are placed relative to the doll in the user's left hand and move with it. Consider arranging books on a bookshelf: the task is easier if the user can see the other books on the bookshelf.

The miniatures in the context are normal dolls: the user can grab one with his right hand and use it to reposition the object it represents. This allows the user to manipulate objects that are occluded or too small on his image plane to select easily. For example, to move a coffee cup that is hidden behind a lamp on a desk, the user creates a doll for the desk and holds it in his left hand; with his right hand he grabs the coffee cup doll from the context and changes its position. The context disappears when the user releases the doll in his left hand or passes it to his right hand.

The context is also useful when the user wants to move an object to a position where there is nothing nearby to serve as a reference object. In this case, the user can create a doll whose context allows at least rough placement. For instance, a user who wants to put a sofa in an empty corner of a room can grab the room as the reference object.

We experimented with a number of rules for choosing dolls to

add to the context. The rule our users preferred for implicitly specifying context is for the system to create dolls for objects within a specified radius and add them to the context. In the current implementation this radius depends on the size of the reference object: dolls are created and added to the context for all objects that are closer to the reference object than twice the longest dimension of the reference object. Thus if a book on a bookshelf is the reference object a few nearby books will be added to the context, but if the bookshelf is the reference object all the books on it will be added.

Another possible rule is to consider semantic relationships, so that objects on a table would be added to the context, but a chair nearby would not. The current system implements semantic relationships through a parent-child hierarchy. When a doll is held in the left hand, the system recursively creates dolls for all the children of the reference object and adds them to the context.



Figure 3: Framing the desired context: the phone and monitor are selected

The user can also explicitly indicate what objects to include in the context. Using both hands the user can frame on the image plane (Figure 3) the objects to be included in the context; the user pinches the thumb and middle finger of his right hand together to start framing, and creates the context by pinching together the thumb and index finger of his left hand. The system adds any object that is completely visible in the frame to the context. Alternately, the user can explicitly specify a radius around a reference object (Figure 4): the system adds any object completely within that radius to the context.



Figure 4: Specifying the radius for the context: the lamp and phone are selected

The user can change the radius of an existing context by pinching the thumb and middle finger of his right hand together and moving his right hand closer to or farther from his left hand. The distance between his hands at which he releases the pinch determines the new radius of the context.

2.4 Animated Objects

Previous techniques make it unnecessarily hard to manipulate or place an object relative to a moving object. Consider the case where a user wants to load a box on a moving truck. The user must either track the truck's motion with his hand, explicitly stop the truck's motion, or match his motion to the truck's before placing the box on the truck. With the Voodoo Dolls technique, the user always works relative to the doll in his left hand, even if the reference object is moving. To put the box on the truck the user creates dolls for the box and the truck, and then puts the box doll on the truck doll without ever changing his own motion or the motion of the truck. However, this technique does not solve the problem that the user must still originally select the truck while it is moving.

2.5 Evaluation

We conducted an exploratory user study with four users to observe [2] people using this technique to arrange furniture in a room. We discovered that, in general, users have little or no difficulty learning to use this technique, and find it a natural way to position objects relative to each other. This study did suggest one significant improvement. When the user is already holding one doll and picks up or creates another, the object represented by the doll in the dominant hand moves instantaneously to the new position determined by the relative positions of the dolls. This caused some confusion if the object was initially in view and suddenly vanished. Instead of instantaneously moving the object, we believe the system should animate the object to its new position over approximately one second.

3 IMPLEMENTATION

We use Fakespace PinchGloves that sense forefinger to thumb contact [12] to detect pinches, and we track the position and orientation of the user's hands by placing a six degree-of-freedom tracker on the index fingers of these gloves. This technique differs from some previous two-handed techniques [8,19] by using gloves, rather than hand held physical *props*. Props provide affordances for interaction, and are most useful when the user works with objects that resemble the props in some way. Because the user can manipulate any object with this technique, there is no "generic prop" that provides affordances for every case.

As with other techniques that use image plane selection, if the world is displayed in stereo the user must designate his dominant eye prior to using the system. The system uses the image displayed to that eye to determine what object the crosshair is over when creating dolls. We currently avoid this difficulty because we render a bi-ocular image that is presented to both eyes. Because dolls created with the right or left hand have different characteristics, the user must also designate whether he is right or left-handed.

4 RELATED WORK

There are a number of interaction techniques for manipulating objects at a distance, but none of them allow the user to work seamlessly at multiple scales.

The Go-Go interaction technique [16] scales the actual motion of the user's hand so objects out of reach can be grabbed. However, the user is still limited by how far he can extend his hand, which makes working at very large scales (e.g. moving buildings around) impossible.

The HOMER technique [1] uses ray casting to select the object, and then moves the user's virtual hand to the position of the

object. The use of raycasting makes it very difficult to select small objects and impossible to select occluded objects. This technique also restricts the manipulation of the object to the physical range of the user's hand. Although the object can be "reeled in," the user must press a button and wait for the object to reach the correct position. This makes it cumbersome to work at large scales.

Stoakley's Worlds in Miniature technique gives the user a miniature hand-held copy of the virtual world. Similar to the Voodoo Dolls technique, manipulating objects in the miniature world affects objects in the larger world. However, shrinking the entire world to a small, hand-held object can make it impossible to select small objects: some objects like pencils or needles may not even be visible. Small movements of objects in the WIM are magnified so that moving a chair a millimeter might move it ten feet in the larger world, making precise placement difficult. Unlike dolls which the user creates dynamically, the user must pre-define the WIMs to create; this fixes the scale the user is working at, rather than allowing this scale to change on the fly. Because the user cannot selectively choose what objects to include in the WIM, clutter may make it difficult to find and select the desired object. The WIM also forces the user to work in a single reference frame: relative to the world. By contrast, the Voodoo Dolls technique allows the user to change the reference frame to suit the task at hand.

With SmartScene [18] the user creates objects by grabbing them from a palette attached to his non-dominant hand. Unlike the voodoo dolls technique, these objects are then placed directly in the scene, not relative to objects held in the non-dominant hand. SmartScene allows users to work at multiple scales, but the user must explicitly change the current scale by stretching or compressing space.

Early work on image plane techniques [15] explored their use for selecting and manipulating objects. Mine [13] proposed shrinking the world after selecting an object with an image plane technique so that the object is within reach to facilitate moving the object in the world. However, this allows only visible objects to be moved, and after the world shrinks the user is still limited by how far he can extend his arm. The user could therefore select a house to move, but would be unable to place it on the other side of town.

Our technique takes advantage of the asymmetric division of labor in positioning objects relative to each other. A number of studies have examined this asymmetry in human bimanual action [3,7,9,10]. In addition to the techniques discussed here, researchers have developed a variety of other desktop and immersive systems that take advantage of the theoretical work in this area [8,17,20].

5 CONCLUSIONS AND FUTURE WORK

The Voodoo Dolls technique allows a user to manipulate objects at a distance by creating miniature copies of objects. These *dolls* have different properties and affect the objects they represent differently when held in the user's right or left hand. This technique overcomes several shortcomings of existing techniques for manipulating objects at a distance. The specific advantages of our technique are:

1. The user can work at multiple scales without explicitly resizing objects or changing modes.
2. Both visible and occluded objects can be manipulated using this technique by either creating the doll directly or grabbing it from the context.
3. Dolls have different effects depending on whether they are

held in the right or left hand. This takes advantage of the asymmetric roles human hands play when manipulating objects.

4. The doll in the left hand provides a stationary frame of reference for the right hand to work in. This simplifies working relative to moving objects.

We plan on extending the Voodoo Dolls technique in two ways. First, we will allow the user create a doll for himself. This will allow the user to move objects to his position if he holds his doll in his left hand, or to move to objects in the scene if he holds his doll in his right hand. Second, we will allow users to create multiple dolls for an object. This will allow users to move an object relative to its initial position by creating two dolls for an object and holding one in each hand.

We believe that other interaction techniques can take advantage of the idea of changing the properties of objects depending on the hand holding them. For example, an object could have control parts that are only visible when the user holds it in his left hand. This allows the user to move or examine the object with his right hand without these parts cluttering his view. If the user wants to manipulate these parts, he transfers the object to his left hand, the controls appear, and he can manipulate them with his right hand.

As another example, consider that the left hand generally makes large, ballistic motions, while the right hand makes fine, precise manipulations [7]. New techniques can take advantage of this by changing the precision or scale of objects' effects. For example, a paintbrush might paint wide, broad strokes when held in the left hand and fine lines when held in the right hand. Similarly, the area erased with the swipe of an eraser could be large if the left hand holds the eraser and small if the right hand holds it. We can even mode traditional desktop objects like dials and sliders in this manner.

6 ACKNOWLEDGEMENTS

We thank Chris Sturgill for his help creating the models and textures for this technique. This technique was implemented in Alice, an authoring tool for building interactive 3D worlds. For more information on Alice, visit <http://www.alice.org>.

This work was sponsored by the Defense Advanced Research Projects Agency (DARPA) and Rome Laboratory, Air Force Material Command, USAF, under agreement number F30602-97-2-0251.

References

- [1] Doug A. Bowman and Larry F. Hodges. An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. *1997 Symposium on Interactive 3D Graphics*, pages 35-38, 1997.
- [2] Frederick P. Brooks. Grasping Reality Through Illusion - Interactive Graphics Serving Science. *Proceedings of CHI '88, 1988*, pages 1-11.
- [3] W. Buxton and B. A. Myers. A Study in Two-Handed Input. *Human Factors in Computing Systems*, pages 321-326, 1986.
- [4] Lawrence D. Cutler, Bernd Frolich, Pat Hanrahan. Two-Handed Direction Manipulation on the Responsive Workbench. *1997 Symposium on Interactive 3D Graphics*, pages 107-114, 1997.

- [5] Scott Fisher, M. McGreevy, J. Humphries, and W. Robinett. Virtual Environment Display System. *1986 Workshop on Interactive 3D Graphics*, pages 77-87, 1986.
- [6] Andrew Forsberg, Kenneth Herndon, and Robert Zeleznik. Aperture Based Selection For Immersive Virtual Environments. *Proceedings of UIST '96*, pages 95-96, 1996.
- [7] Yves Guiard. Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *The Journal of Motor Behaviors*, 19(4):486-517, 1987.
- [8] Ken Hinckley, Randy Pausch, John C. Goble, and Neal F. Kassel. Passive Real Worlds Props for Neurosurgical Visualization. *Proceedings of CHI*, pages 452-458, 1994.
- [9] Ken Hinckley, Randy Pausch, and Dennis Proffitt. Attention and Visual Feedback: The Bimanual Frame of Reference. *1997 Symposium on Interactive 3D Graphics*, pages 121-126, 1997.
- [10] Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassel. Cooperative Bimanual Action. *Proceedings of CHI*, pages 27-34, 1997.
- [11] P. Kabbash, W. Buxton, and A. Seken. Two-Handed Input in a Compound Task. *Proceedings of CHI*, pages 417-423, 1994.
- [12] D. P. Mapes and J. M. Moshell. A Two-Handed Interface for Object Manipulation in Virtual Environments. *Presence*, 4(4): 403-416, 1995.
- [13] Mark R. Mine, Frederick P. Brooks, Jr., and Carlo H. Sequin. Moving Objects in Space: Exploiting Proprioception in Virtual Environment Interaction. *ACM SIGGRAPH '97 Conference Proceedings*, pages 19-26, 1997.
- [14] Randy Pausch, Tommy Burnette, Dan Brockway, Michael E. Weiblen. Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniatures, *ACM SIGGRAPH '95 Conference Proceedings*, 1995.
- [15] Jeffrey Pierce, Andrew Forsberg, Matthew Conway, Seung Hong, Robert Zeleznik, and Mark Mine. Image Plane Interaction Techniques in 3D Immersive Environments. *1997 Symposium on Interactive 3D Graphics*, pages 39-44, 1997.
- [16] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR. *Proceedings of UIST 1996*, pages 79-80, 1996.
- [17] Chris Shaw and Mark Green. Thred: A Two-Handed Design System. *Multimedia Systems Journal*, 3(6), 1995
- [18] SmartScene™ is a product of MultiGen Inc. More information on SmartScene™ is available from MultiGen's website at <http://www.multigen.com>.
- [19] Richard Stoakley, Matthew Conway, Randy Pausch. Virtual Reality on a WIM: Interactive Worlds in Miniature. *ACM SIGCHI '95 Proceedings*, 1995.
- [20] Robert C. Zeleznik, Andrew S. Forsberg, and Paul S. Strauss. Two Pointer Input for 3D Interaction. *1997 Symposium on Interactive 3D Graphics*, pages 115-120, 1997.