SnowGlobe: A Spherical Fish-Tank VR Display

John Bolton
Human Media Lab
Queen's University
Kingston ON, Canada
bolton@cs.queensu.ca

Kibum Kim
Human Media Lab
Queen's University
Kingston ON, Canada
kibum@cs.queensu.ca

Roel Vertegaal
Human Media Lab
Queen's University
Kingston ON, Canada
roel@cs.queensu.ca

Copyright is held by the author/owner(s).
CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract

In this paper, we present a spherical display with Fish-Tank VR as a means for interacting with three-dimensional objects. We implemented the spherical display by reflecting a projected image off a hemispherical mirror, allowing for a seamless curvilinear display surface. Diffuse illumination is used for detecting touch points on the sphere. The user's head position and the position of the sphere are also tracked using a Vicon motion capture device. Users can perform multi-touch gestures to interact with 3D content on the spherical display. Our system relies on the metaphor of a snow globe. Users can walk around a display while maintaining motion parallax corrected viewpoints of the object on the display. They can interact with the 3D object using multitouch interaction techniques, allowing for rotating and scaling of the 3D model on the display.

Keywords

3D Interaction, Motion Capture, Tangible UI, Organic UI, Haptic Interfaces.

ACM Classification Keywords

H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms

Design, Human Factors, Performance.

Introduction

With the advent of new thin-film display technologies, such as Organic LEDs and E-Ink displays, it has become conceivable that one day, any surface can be a display. This means form factors will likely not be limited to flat surfaces. Spherical form factors, for example, have the advantage of allowing users to access 3D data from any viewpoint. They have been prevalent in Earth globes for this reason. However, spherical displays are currently limited to the display of spherical information, such as maps. Due to the curvature of a spherical display, it is difficult to display flat information or correctly represent a 3D scene.

In this paper, we discuss the design of a 3D object viewer that uses touch input and head tracking to mimic volumetric projection, where a 3D object appears to be sitting inside the display volume. We have created a set of simple gestural techniques for interacting with 3D objects on the display. Gestured input is tracked using an optical multi-touch system contained within the sphere and the user's head position is tracked using a Vicon motion capture system. This allows our system to preserve motion parallax and present the correct 3D perspective of 3D objects when walking around the display.

With this project, we seek to examine the feasibility of using motion-parallax 3D projection on a sphere to display 3D objects as if they were volumetric. According to past thought experiments [2], this is categorized as a "nonplanar 2D to 3D mapping". The user can visualize the object they are viewing as being positioned at the center of the sphere, with interaction on the spherical surface mapped to the object's rotational axes. This provides a more natural feel than traditional planar alternatives. Additionally, we allowed for scaling through the interaction with the display, allowing users to focus on specific portions of the model. The vocabulary of gestural interactions was designed to be very simple, mapping directly to how a spherical surface containing a physical object would behave in the real world. We expect that this minimal interface creates a compelling and much more affordable and high resolution alternative to a true volumetric display.

Related work

With recent theoretical [2] and practical [1] work in real-time motion-capture gestures and 3-dimensional displays, research on displays of any shape or form has become more prevalent.

In the past, the editing of 3D objects by direct gestural interaction has been studied for regular planar interfaces [6] and there has been work with motion-captured gestures to edit 3D objects displayed on a separated flat surface [7] or on a 3D volumetric Actuality's Perspecta display [1]. However, these projects did not explore the use of 3D projection and motion parallax compensation on a spherical surface to view 3D objects. Sheng et al. [7] examined the use of a physical object, such as a sponge, as a proxy for the user to hold and deform while gestures were detected via motion-capture. Here, the 3D object being edited was displayed on a separate screen.

Stavness et al. [8] developed pCubee, a perspective handheld display. Five small LCD panels were arranged to form the sides of a cube. Head-coupled perspective rendering was used to simulate the feel of real objects located inside the cube. Additionally, a real-time physics engine was implemented allowing for manipulation of the displayed objects. To evaluate their system, users were required to complete a 3D tree-tracing task. They found that bimanual interaction techniques were preferred to other interaction conditions.

Grossman et al. [4] developed a 3D geometric model building application to demonstrate multi-finger gestural interaction on a hemispherical volumetric display. The user's fingers were tracked using a Vicon motion tracking system. Interaction techniques were designed to make use of the unique features of volumetric displays, specifically, the 360 degree nature of the viewing volume. Their set of interaction techniques consisted of *SurfaceBrowser* for file management, model transformations, and techniques to combine models to create scenes. *SurfaceBrowser* could be rotated by scrubbing the non-dominant hand's index finger along the display surface, bringing objects into view that were originally hidden. Alternatively, the user could walk around the display to view this information. Model transformations consisted of rotation, translation, and scaling. They made the observation that the 360-degree visibility of information when using a volumetric display makes it useful for exploring collaborative multi-user interaction. Due to the unique properties of a 360 degree display, considerations of how users share space must be examined. Additionally, it was found that it was useful for the *SurfaceBrowser* to show information on the inner surface of the display, in effect, reducing it from a 3D volumetric display to a 2D hemispherical display.

Figure 1. SnowGlobe setup, showing our spherical display.

SnowGlobe Implementation

Our spherical display surface consists of a 36″ diameter hollow sphere made of acrylic, with a 19″ hole cut in the bottom (see Figure 1). The acrylic was sandblasted to create a diffuser on the inside, allowing it to act as a projection surface. This also allows it to act as a diffuse illumination surface for infrared light, enabling multi-touch detection. A DepthQ 3D projector is mounted at the bottom of a base, a $35 oil drum commonly known as a burning barrel, using a projector mount inside of the barrel. The projector is aimed upwards in order to project a 1024x768 image off an 18″ hemispherical mirror mounted inside the top part of the sphere. This mirror projects light to the rest of the surface of the display. Infrared illuminators are pointed at the inside of the acrylic sphere and arranged to spread light evenly across the surface. Two SONY EyeToy cameras with wide-angle lenses function as infrared cameras. These are placed inside the sphere pointing towards

Figure 2. Rotation gesture. To rotate a model a user, uses a single touch. The axis of rotation is determined by whether the user moves their finger horizontally across the sphere's surface or vertically. This analogous to spinning the entire volume.

Figure 3. Scaling gesture. To scale a model, a user uses a two finger pinch. As the distance between the fingers increases, the models is enlarged as the distance decreases, the model is shrunk.

each hemisphere. SnowGlobe uses a diffuse illumination approach to register multi-touch points from the fingers of a user on the surface of the sphere.

Vicon motion-capture is used [5] to determine the location of the sphere and the user's head position. The Vicon consists of 8 cameras, which are placed around the sphere. To see 3D content on the sphere, users wear a pair of shutter glasses with three reflective infrared markers, one on each side, and one on the bridge above the nose. To track the location of the sphere, we affixed four reflective infrared markers in an imaginary square around the north pole of the sphere. An overview of the setup is shown in Figure 1.

Vicon tracking output is parsed in real-time by a Quartz Composer patch. Changes in the angles between the user and the center of the sphere are used for motion parallax compensation, and rotate the current model and the projected image as the user moves around the display. This allows the model to always be at what appears to be a static position on the display.

Figure 4. Model switching gesture. To change models, a user uses a full handed swipe along the sphere. This is designed to represent the act of pushing the model out of the display.

The projected image is rendered by applying a concave bulge to the user's current view and then mapping this image to a hemispherical surface. By projecting this image onto our hemispherical mirror inside the sphere we are able to project a 3D object with minimal distortion.

Our system can import 3D objects into the Quartz Composer application as Collada files. These models can be found in the Google Sketchup 3D Warehouse. Models are scaled to a unit cube on import, allowing them to be shown at an optimal size for the display surface.

Interaction techniques

Gestures are determined by tracking touches along the surface of the sphere and are subsequently mapped to actions on the 3D object inside the sphere. We designed three basic gestural interactions for use with the sphere: Rotation (Figure 2), Scaling (Figure 3), and Swiping (Figure 4).

Rotation. Rotation of the model is limited to the x and y axes. By touching the sphere at one point and moving along an imaginary line of latitude the model will be rotated along the y axis. By placing one touch on the surface and moving along a longitude, the model will be rotated along its x axis. This gesture is based on the concept of the user spinning the sphere as if it was a free-spinning globe.

Scaling. Scaling is performed as a bimanual gesture. By placing two touches on the spherical screen the user can uniformly scale the model, using a pinching gesture. If the two touches move closer together the model is shrunk while movement of the touches away from each other will enlarge the model. This gesture is common in multi-touch interactions and should be easily recognizable to most users.

Swiping. A user can change the current model by using a full handed scroll gesture of sufficient speed. By placing a hand on the display surface and quickly moving it along a latitude in a clockwise manner the user can navigate through available models. If the scroll is in the opposite direction the user can go backwards through the list. This gesture was designed to mimic the idea of pushing an object out of the display volume.

Discussion

We believe our 3D viewer may help with the visualization of 3D objects, which can be difficult on flat panel screens. Additionally, we believe that we can duplicate the advantages of volumetric displays in a more high resolution and much less expensive ($4000) system. The utility of our system could be improved by developing additional applications. Our model viewer does not support the deformation or manipulation of objects. We are considering additional use cases, such as virtual clay modeling.

Currently, our system only supports one user but this limitation is easily overcome by adding additional head tracking. By allowing two users to use the display concurrently, we can present two completely different views of an object making use of a spherical display's characteristic of only providing a single viewable hemisphere to a user at a time. This would allow for private collaboration on a single display surface while maintaining a sense of two users working on a single task, which is currently not possible on a volumetric display.

Conclusions

In this paper, we presented SnowGlobe, a multi-touch sensitive spherical fish-tank VR system that can be used to manipulate 3D objects. We implemented SnowGlobe using a 3D projector and Vicon head tracker. By tracking the user's head position, we are able to preserve motion parallax of 3D objects. Objects rotate as the user walks around the display, making it appear as if he is walking around them. Our gesture set allows users to rotate and scale objects using multitouch on the spherical display surface. Initial experiences suggest this display gives the illusion that the object on the display is suspended within the 3D space of the globe.

Acknowledgements

This work was funded by an Ontario Research Fund, with support by Autodesk Research and fuseproject.

References

[1] Actuality's Perspecta display, http://www.actuality-medical.com/

[2] Balakrishnan, R., Fitzmaurice, G., & Kurtenbach, G. User interfaces for Volumetric Displays*. IEEE Computer 34,* 3 (2001), 37-45.

[3] Benko, H., Wilson, A. D., and Balakrishnan, R. Sphere: multi-touch interactions on a spherical display. In *Proc. UIST 2008*, ACM Press (2008), 77-86.

[4] Grossman, T., Wigdor, D., and Balakrishnan, R. Multi-finger gestural interaction with 3d volumetric displays. In *Proc. UIST 2004*, ACM Press (2004), 61-70.

[5] Holman, D., Vertegaal, R., Altosaar, M. Paper Windows: Interaction Techniques for Digital Paper. In *Proc. CHI 2005,* ACM Press (2005), 591-599.

[6] Nishino, H., Utsumiya, K. and Korida,K. 3D Object Modeling Using Spatial and Pictographic Gestures. In *Proc. VRST 1998*, ACM Press (1998), 51-58.

[7] Sheng, J., Balakrishnan, R. and Singh, K. An Interface for Virtual 3D Sculpting via Physical Proxy. In *Proc. Graphite 2006*, ACM Press (2006), 213-220.

[8] Stavness, I., Lam, B. and Fels, S. pCubee: a perspective-corrected handheld cubic display. In *Proc. CHI 2010*. ACM Press (2010), 1381-1390.