VORTEX: Design and Implementation of an Interactive Volumetric Display

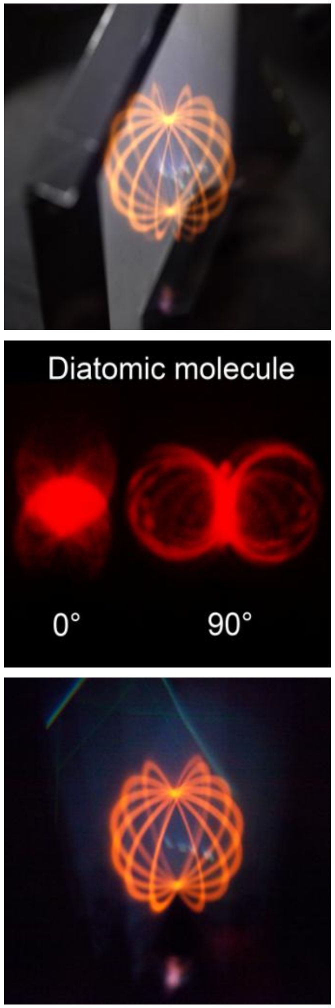

Figure 1. 3D Volume renders as seen on VORTEX.

Abhijit Karnik
Dept. of Computer Science
University of Bristol, UK
karnik@cs.bris.ac.uk

Archie Henderson
Dept. of Mechanical Engg.
University of Bristol, UK

Andrew Dean
Dept. of Mechanical Engg.
University of Bristol, UK

Howard Pang
Dept. of Mechanical Engg.
University of Bristol, UK

Thomas Campbell
Dept. of Mechanical Engg.
University of Bristol, UK

Satoshi Sakurai
Human Interface Engineering Lab
Osaka University, Japan
sakurai.satoshi@ist.osaka-u.ac.jp

Guido Herrmann
Dept. of Mechanical Engg.
University of Bristol, UK
g.herrmann@bristol.ac.uk

Shahram Izadi
Microsoft Research
Cambridge, UK
shahrami@microsoft.com

Yoshifumi Kitamura
Interactive Content Design Lab
Tohoku University, Japan
kitamura@riec.tohoku.ac.jp

Sriram Subramanian
Dept. of Computer Science
University of Bristol, UK
sriram@cs.bris.ac.uk

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

Abstract

True 3D display systems like volumetric displays allow generation of autostereoscopic, multi-view 3D content that has real physical dimensions. However their uptake as a research tool within the HCI community is limited largely due to difficulties in buying or building such displays. The choice of commercially available systems is limited and constrains the flexibility of their use in terms of interaction capabilities, display features and integration with multi-display environments (MDEs). In this paper we describe the steps involved in creating custom volumetric display from easily available components. By building a touch-enabled volumetric display we walk-through the steps involved in the process. This will enable us to explore various interactive systems, associated techniques and challenges related to integration of the device into a MDE.

Keywords

Volumetric display, true 3D, isotropic displays

ACM Classification Keywords

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

General Terms

Design, Human Factors

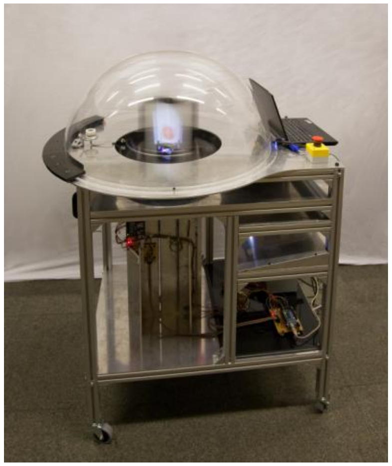

Figure 2. VORTEX. ‘Volumetric 3D Explorer’ in operation.

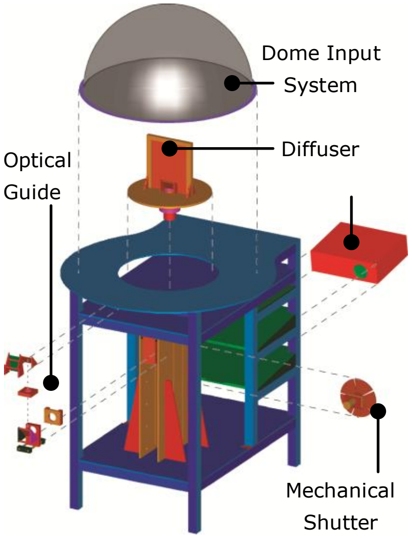

Figure 3. VORTEX exploded view.

Introduction

Volumetric displays generate views in which the displayed objects have real spatial dimensions. Most examples of volumetric displays are one-off research prototypes such as the DepthCube [5], Laser Induced Damage glass [4], and Lightfield 360° [3]. The only commercially available volumetric display used to be from Actuality System [1]. These mostly research prototypes demonstrate that the graphics and the hardware research community have the knowledge to build such displays. An often lamented benefit of these displays is the user’s ability to better perceive the 3D data and support more intuitive interaction possibilities.

However due to limited access to such displays within the HCI community researchers have not really examined the interaction issues around such displays (only notable exception is the work from University of Toronto such as [2]) or examined how these displays can be integrated as part of an end-users task flow. A custom built MDE consisting of an interactive volumetric device and other planar devices can provide a content-rich interactive multi-user environment. The implementation of such a system is non-trivial. We need to identify the issues related to interaction techniques and to the arrangement of a virtual environment across dimensionally disjoint display systems.

However, the very first issue is that volumetric devices are not readily available to the HCI community. Democratization of access to volumetric displays can spawn a research direction in a similar way to multi-touch tables. The ready availability of know-how to create multi-touch tables has led to a strong research and DIY community around multi-touch tables. We feel that the moment is right for these communities to explore the use of volumetric displays alongside multi-touch tables.

Hence we chose to channel our efforts in building VORTEX (Figure 2), an isotropic swept volumetric display using off-the-shelf hardware. The device uses OpenGL to render 2 volumes per second with 420x420 voxels having 3° angular width using a 120Hz NEC NP410 projector in 24-bit color.

In this paper we present the design and features of VORTEX. We explain the steps involved and expose potential pit-falls in designing custom touch-enabled volumetric displays. Finally we identify the open research challenges in integrating volumetric displays with other planar devices to create hybrid MDEs.

Theory of Operation

To generate a 3D object in the volumetric space the volumetric points/pixels (voxels) of interest are converted into point sources of light. The two well-known methods of manipulating voxels are static-volumes and swept-volumes. Static volumes utilize the bulk properties of materials to trigger the voxels [5, 4]. Swept-volumes rely on the presence of a moving diffuser to periodically ‘sweep’ each point in space and portray the voxel information relevant to its spatial position inside the volume [1, 3]. Holographic systems are currently incapable of generating dynamic interferograms and hence not considered as volumetric displays.

Isotropic displays are a sub-class of swept-volume displays. These displays render the volume assuming that each point in the volumetric space is visible to

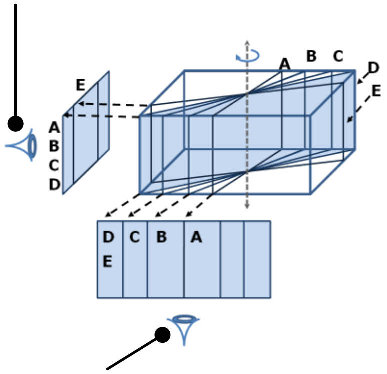

User perceives the edges from frame A to D as same size but one edge recedes while other advances

User perceives the edges from frame D and E as same size but one edge recedes while other advances

Figure 4. Volume Slicer visualization. A rectangular cuboid is rendered as a series of rectangles (A, B, C, D & E)

every viewer in any direction. The diffuser is designed to spread the incident light rays as uniformly as possible. This creates the 'translucent skin' effect which is interpreted as a constraint of these displays. It can however be a useful feature in certain scenarios requiring cut-away views of complex structures. The other feature is that the display renders a *360^o accessible, view-angle independent volume*. Thus it is not necessary to track the users. The perception of depth is due to actual spatial orientation of the projected objects and with natural parallax. These features make it an ideal candidate for implementation.

VORTEX, being isotropic, renders the 3D universe in a piece-by-piece basis only considering the angle of the display screen. This can be visualized as a slicing plane rotating around a vector parallel to the plane itself passing through the objects of the universe to be rendered. At any point in time, the slice data consists of the edges of intersection of the objects lying in the same position as the plane. The edge information can be used to generate the surface as standalone representation. Depending on the application, it can also be used to fill the area bounded by the edges to mimic a physical slicer that generates cross-sections of the object as it moves through it.

In the standalone form, a rectangular cuboid primitive is rendered as a series of radially aligned rectangles that change in size. The perceived height and width of the rectangles remains the same for a user whose line of sight is along the progression of vertical edges. The effect is maintained since one vertical edge of the rectangle recedes from the user while the other one advances (refer Figure 4). A sphere is rendered as a set of circles depending upon its position. If the sphere is at the origin, the circles are arranged like the longitudes on its surface. For an off-centric sphere, the circles will not intersect and will vary in size. All primitives can be visualized in a similar manner. See Figure 1 for sample renders on VORTEX.

Implementation Details

The key components of VORTEX are the projection system, the optical guide, the diffuser, mechanical shutter, volume render engine and input system (Figure 3). These are described next:

Projection System and Optical Guide:
A LCD projector is used to generate the volume image on the screen. Since a projector is designed to project a large image at a distance, the image has to be shrunk and focused on a smaller screen placed closer than the minimum distance specified for the projector. We used an optical guide path consisting of two lens stages. First stage was a 50mm f/1.8 Olympus OM Zuiko MC lens in place of plano-convex lenses which are prone to chromatic aberration. Stage two was a plano-concave diverging lens. Mirrors were used in between to bend the light path.

Diffuser and Mechanical Shutter:
VORTEX sweeps 2 volumes a second. Since one volume is swept in 180° rotation of the diffuser, the diffuser only needs to rotate once every second (1 Hz). A translucent screen (50% Transmission – 50% Reflection) allows a viewer to continue seeing projected contents throughout the rotation. However 1 Hz is slow enough for a user to just track the diffuser. At 15 Hz, it is nearly impossible to do so and anything projected on the diffuser appears to float in mid-air for which reason this higher frequency of 15 Hz is practically used. The

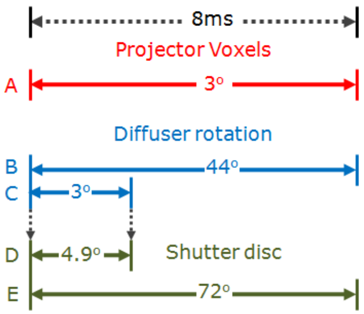

Figure 5. Decoupling the projector refresh rate from the diffuser rotation rate.
A: Actual angular width of voxels to be displayed (takes 8ms @ 120fps)
B: Angle swept by diffuser at 15 Hz
C: Actual swath of volume to be lit.
D: Open slit on shutter disc allowing light through to diffuser
E: Angle rotated by shutter disc in 8ms. (Disc has 5 slits)

Figure 6. Input System.
FTIR on the dome surface can be used as a source of interaction inputs to manipulate the volume.

user's focus shifts from the motion of the diffuser to what is being displayed on the diffuser.

However at 120 frames per second (fps), the projector takes 8.3 milliseconds between changeovers from one frame to another. During this time, the screen will have moved by (15*360*8.3/1000=) 44.82°. The result is a motion artifact that allows for the data meant for 3° to be shown for over ~44° which leads to blurring. VORTEX overcomes this limitation by introducing a shutter-disc mechanism with slits allowing light to pass through. The timing of this system is shown in Figure 5. Due to this the render engine renders non-consecutive slices (i.e. 1°, 15°, 30° instead of 1°, 2°, 3°).

A brushless DC motor drives the diffuser while a brushed DC motor drives the slit assembly. An Arduino Mega carries out the speed control functions by outputting commands to the motor driver cards. The Arduino Mega is controlled at a higher level through a USB-Serial interface. The device structure was made up of Bosch aluminum framing and parts laser cut from aluminum sheets.

Volume Render Engine:

The position of the diffuser determines how the volume is sliced (as described previously). We used a quadrature optical encoder connected to a Phidget 1057 High Speed Encoder (0.8° resolution) to get this position. The default refresh rate of the Phidget's USB interface is 125/second. This rate is too slow for our application where 8 milliseconds can push the diffuser out of synch by ~44°. As a solution to this issue, we introduced a test-mode driver designed to deliver a 1000/second refresh rate. An OpenGL program renders the clipped slice based on the position of the screen.

Input System:

Since one of the reactions to a volumetric display is the attempt to 'poke' or touch it, (and which can lead to serious physical injury), the protective dome separates the user from the fast spinning diffuser. This protective dome also serves as an input surface for VORTEX. By injecting IR into dome flange, frustrated total internal reflection (FTIR) was achieved. When a user touches the dome, a camera placed below captures the touch input (see Figure 6) and converts it into interaction gestures for the volume. This is unique and has not been previously reported for curved surfaces like the dome that we are using.

Factors Affecting Display Parameters

The display parameters of a volumetric display viz. voxel count (10,584,000), voxel angular width (3°), color depth (24-bit) and volume refresh rate (2 volumes/s) are interrelated and determined by a set of system dependent factors. The resolution of a volumetric display is the pixel resolution of a single 2D frame (slice) it can render, scaled by the number of slices displayed in a single second. Since a projector is used to generate the slice, the following projector parameters dictate the voxels count.

Bandwidth:

For a volume made up of N slices with n bit color depth, each with a width W and height H, redrawn every R times/second requires the data links to transfer data for n * N * W * H * R bits/s (2.2 Gbps). While specifications like DVI-D can handle the data transfer rate, commercially available systems rarely require such data rates. Consequently the hardware has limited data sink capabilities.

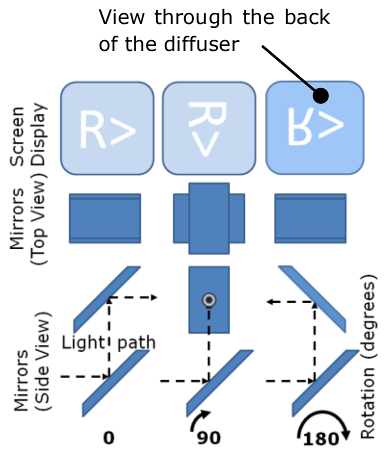

Figure 7. Optical Correction.
The effect of optical artifacts when they are not corrected by the render engine.

Vertical Refresh Rate:
The vertical refresh rate (VSync) of a projector represents the practical maximum for fps delivered by the projector. Since each frame corresponds to a slice, the maximum slices/second is limited accordingly. Hence VSync also limits the volume refresh rate which is a function of slices/second.

Color Depth:
With the limitations mentioned for frame rates, a middle ground is needed to achieve higher frame rates. Options like pre-loaded renders [3] require specialized development boards. The color depth provides trade-off for increasing the number of slices being displayed by the projector if the projector allows 16, 8, 4 or 1-bit color depth modes. Scaling down the color depth allows us to scale up the voxel rate by the same factor. VORTEX currently retains full color depth and settles for the highest projector capability of 120 fps.

Optical Corrections:

The projector system is kept below the actual display screen. An optical lens-mirror assembly guides the light to the screen. As the platform rotates away from the start position, the two mirrors are axially aligned at an angle. The image formed on the diffuser is rotated in the opposite direction as that of the rotation of the platform. To compensate for this, the perspective camera in the rendering software rotates its Up-vector at an angle equal to the rotation of the platform.

With the above correction, when the diffuser rotates 180°, the user sees the image that has passed through the diffuser instead of what was reflected by the diffuser. The user sees a mirrored flip of the desired image. To remove this artifact, we revolve the camera perspective around the origin of the universe along with the up-vector rotation. This ensures the relative orientation of the slice is the mirror inversion of the object thus cancelling out the artifact. These artifacts are demonstrated in Figure 7.

Inter-frame Gap – An Error Source:

At 120 fps, a new frame should be displayed at every 8ms. However the inter-frame gaps can vary significantly (between 6ms to 12ms but with average value of 8ms) as seen in Figure 8. This was seen with different graphics cards and machine configurations. These variations are detrimental to the timing precision required to show a slice at an exact location on the volume. For every 1 ms of variation to the frame start time, the image is wrongly positioned by 5°. Even with the mechanical shutter the delayed slice gets displayed in a wrong location. This issue is not apparent with 2D displays as the frames are not spatially separated. Hence the issue is not well documented. We used an NVidia Quadro FX 4800 card along with the GSync-2 card for this. Only by synchronizing the frame timing with a 120 Hz external Genlock signal supplied to the GSync-2 card, the inter-frame gap was restored to 8ms.

Future Work (Device Specific)

VORTEX opens up additional research challenges that need to be addressed. These are:

- Higher volume refresh rate (VRR): Work needs to be done to see if a multi-projector setup can scale VRR. This can improve the realism of the display. Challenges with this approach would be the precise matching of the light paths and synchronization of the shutter mechanisms.

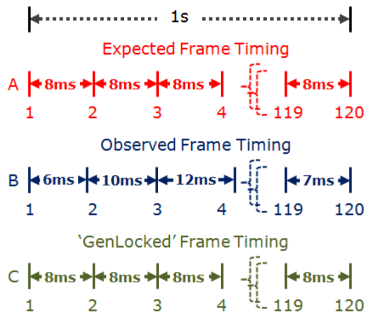

Figure 8. Inter-frame Gap Issue. The inter-frame gaps as seen by the projector.
A: Expected frame timing necessary for correct projection of slices.
B: Observed frame timing with standalone graphics card.
C: Observed frame timing with GSycn-2 card (120 Hz external GenLock signal source)

- Dynamic diffuser: Explore if a non-monolithic, non-interacting diffuser (e.g. air) is technically feasible. This will allow the user to safely reach into the volume.
- Interaction techniques: The input capabilities of the dome need to be explored further so as to transform the curved input space into a volume transformation space. While traditional input systems can work, the dome provides the *'walk-up and play'* feature of multitouch tables.

Future Work (Hybrid MDE Approach)

- Building a hybrid MDE: Currently there is a gap in literature related to integration of VORTEX-like devices with a standard planar MDE. The integration process is non-trivial as the input space and the display space of the systems are dimensionally disjoint.
- Object behaviour: Examine how an object is represented on a single device and during inter-device transfer situations.
- Interactions: Leverage the interactive capabilities of VORTEX in conjunction with the rest of the MDE and explore interaction techniques relevant to such a system.
- Real world deployment: There is no data available on the performance of a hybrid MDE in a real world scenario. This requires us to implement an end-user task and evaluate the performance of the system as a whole.

These challenges are the focus of our continuing work in this area.

Summary

We built VORTEX to demonstrate the feasibility of an ultra-low cost swept-volumetric device using off-the-shelf components. We also implemented an FTIR based input interface for VORTEX. We discuss the design constraints affecting VORTEX as well as other swept volumetric devices following a projector based approach. This work serves as a primer for anyone wishing to build a similar interactive volumetric display device. We highlight the observations made during the process of implementing VORTEX. Some of these observations are generic to all forms of volumetric displays irrespective of their implementation schemes. We also present potential challenges for future work on such volumetric systems.

Acknowledgements

The authors thank Timo Kunkel, Jason Alexander, Izdihar Jamil, Benjamin Long and other members of the Bristol Interaction & Graphics Group, University of Bristol for their feedback throughout the genesis of VORTEX. This work is supported by Microsoft Research through its PhD Scholarship Programme.

References

[1] Favalora, G., Napoli, J., Hall, D., Dorval, R., Giovinco, M., Richmond, M. and Chun, W. S. 100-million-voxel volumetric display. *SPIE* 4712: 300-312. (2002)
[2] Grossman, T. and Balakrishnan, R. The design and evaluation of selection techniques for 3D volumetric displays. *UIST '06*, 3-12. (2006)
[3] Jones, A., McDowall, I., Yamada, H., Bolas, M. and Debevec, P. An interactive 360° light field display. *SIGGRAPH '07*, 13. (2007)
[4] Nayar, S. K. and Anand, V. N. 3D Display Using Passive Optical Scatterers. *Computer* 40(7): 54-63. (2007)
[5] Sullivan, A. DepthCube solid-state 3D volumetric display. *SPIE Electronic Imaging '04*, 279-284. (2004)