BalloonProbe: Reducing Occlusion in 3D using Interactive Space Distortion

Niklas Elmqvist
Department of Computer Science
Chalmers University of Technology
SE-412 96 Göteborg, Sweden
elm@cs.chalmers.se

ABSTRACT

Using a 3D virtual environment for information visualization is a promising approach, but can in many cases be plagued by a phenomenon of literally not being able to see the forest for the trees. Some parts of the 3D visualization will inevitably occlude other parts, leading both to loss of efficiency and, more seriously, correctness; users may have to change their viewpoint in a non-trivial way to be able to access hidden objects, and, worse, they may not even discover some of the objects in the visualization due to this inter-object occlusion. In this paper, we present a space distortion interaction technique called the BalloonProbe which, on the user's command, inflates a spherical force field that repels objects around the 3D cursor to the surface of the sphere, separating occluding objects from each other. Inflating and deflating the sphere is performed through smooth animation, ghosted traces showing the displacement of each repelled object. Our prototype implementation uses a 3D cursor for positioning as well as for inflating and deflating the force field "balloon". Informal testing suggests that the BalloonProbe is a powerful way of giving users interactive control over occlusion in 3D visualizations.

Categories and Subject Descriptors

H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*artificial, augmented, and virtual realities*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*virtual reality*

General Terms

Human Factors, Design

Keywords

occlusion reduction, 3D space distortion, interaction technique

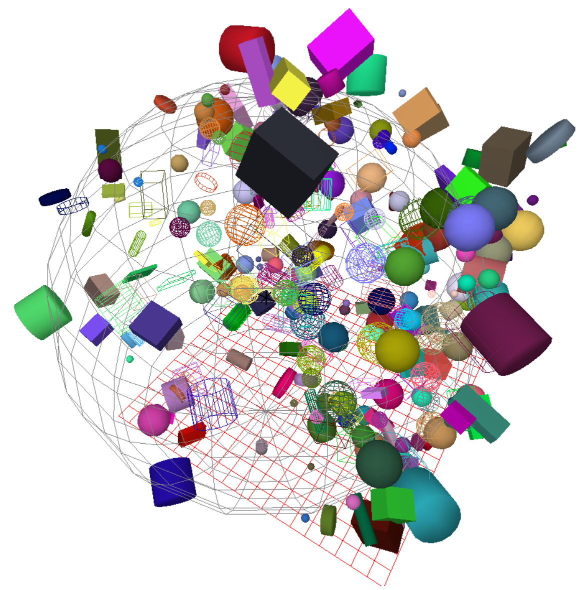

Figure 1: The BalloonProbe interaction technique in action on a moderately complex 3D scene.

1. INTRODUCTION

There is strong potential for making use of advances in the field of information visualization inside 3D virtual environments, where users can effortlessly overview and interact with the visualized data using natural interaction methods and multimodal techniques. However, the immersiveness of a virtual environment may also be a hindrance to the effectiveness of using the visualization; objects in a 3D environment will inevitably occlude other objects, partially or totally, leading to a phenomenon of literally not being able to see the forest for the trees. This inter-object occlusion problem applies to all 3D visualization, and can either result in reduced performance by forcing the user to change the viewpoint to access hidden objects, or even in reduced correctness due to some objects remaining undiscovered by the user.

The intuitive approach employed in this paper is to equip the user with a portable 3D force field that can be inflated and deflated to separate occluding objects, much akin to a balloon that pushes nearby objects away from each other as it is inflated. This interaction technique, accordingly called a BalloonProbe (see Figure 1 for an example screenshot), is freely controlled by the user in 3D space and is smoothly animated when it is inflated and deflated. Ghosted traces of the original position of the displaced objects are drawn to help the user maintain a cognitive overview of the original 3D scene. The displaced objects can optionally be subject to collision detection to avoid overlaps on the surface of the balloon, but this is really only necessary for special cases when an object is fully contained by another.

We have implemented a prototype of the BalloonProbe technique on a standard PC using a mouse and keyboard for input, as well as a low-budget 6DOF motion tracker. Informal testing using a few test subjects indicate that the BalloonProbe is a very useful tool for disambiguating complex 3D visualizations with a high degree of inter-object occlusion.

2. RELATED WORK

Three-dimensional computer graphics have great potential for visualization applications, but also carry a number of intrinsic problems that threatens its usefulness for information visualization, inter-object occlusion being only one of them. Traditionally, work in this field has been concentrated on navigation and orientation within 3D environments, and many papers exist on this subject.

Some work dealing more directly with object discovery and access in complex 3D environments has also been conducted; the Worlds-in-Miniature technique [7] uses a miniature 3D map of the environment to support both discovery and access, worldlets [1] provide both global overview maps as well as local views optimized for detail, and bird's eye views [2] combine overview and detail views of the world.

The BalloonProbe technique disambiguates occluded objects using a form of 3D space distortion, a general approach in information visualization that was introduced by Furnas in his work on fisheye views and focus+context techniques [3]. Other papers where space is distorted to favorably present certain objects over others include the Perspective Wall [5] for visualizing linear information on a 3D surface, the Table Lens [6] for spreadsheet-like tabular visualization, and the Hyperbolic Tree Browser [4] for graphically representing large hierarchies using hyperbolic geometry. In particular, the EdgeLens [8] method bears many similarities to the BalloonProbe, especially its bubble variant, but is intended for selective reduction of edge congestion in 2D graphs.

3. THE OCCLUSION PROBLEM

We represent the 3D world U by a Cartesian space $(x, y, z) \in \mathbb{R}^3$. Objects in the set O are volumes within U (i.e. subsets of U). A ray r is said to intersect the object o if the ray intersects any part of o. A ray r is **blocked** by an object o if r intersects o. An object o is said to be **occluded** (or fully occluded) from a viewpoint v if all rays r between v and o are blocked. Analogously, an object o is said to be **visible** from a viewpoint v if there exists a ray r between v and o such that r is blocked by no object. An object o is said to be **partially occluded** from viewpoint v if o is visible, but there exists a ray r between v and o such that r is blocked by another object.

Given the above definitions, we can categorize the general occlusion problem as consisting of two main parts:

i) *object discovery* – finding all objects $o \in U$ in the environment, and

ii) *object access* – retrieving graphically encoded information associated with each object.

Object discovery efficiency is severely hampered by the existence of fully occluded objects, whereas object access also suffers for partially occluded objects. Both issues will affect the efficiency and correctness of users solving tasks using a visualization, but clearly, object discovery is the most serious of the two; if the user is unaware of the existence of an object, she will have no motivation to look for it and access never becomes an issue. The interaction technique presented here will help both subproblems, both by separating occluded objects so that the user can access them, and by making previously unknown objects visible so that they can be discovered.

The most common way to overcome partial or total occlusion is to change the position and orientation of the viewpoint. This is akin to the real world, where occluded objects are discovered and accessed by walking around in the 3D environment. In this paper, we instead distort the 3D space as if the user had a inflatable balloon that can push objects away from each other to reduce the occlusion.

4. INTERACTIVE 3D SPACE DISTORTION

The basic idea behind the BalloonProbe interaction technique is to distort the space locally in areas of high object congestion, introducing a spherical "force field" that repels objects in the visualization to help the user disambiguate between them. This force field is similar to a balloon in that it can be inflated and deflated on the user's command, and the transition from zero to full size and back is performed through a smooth animation. Objects that are in the way of the expanding balloon are pushed away accordingly, and then resume their original position as the balloon is deflated. The effect is that objects are projected onto the surface of the spherical balloon where they can be easily accessed by orbiting the camera around the balloon centerpoint.

See Figure 2 for a schematical 2D overview of the balloon distortion (this overview trivially generalizes to 3D). The balloon itself is centered around the 3D cursor, controlled by the user, who can not only toggle the inflation/deflation of the balloon, but also modify its full size depending on the situation. This gives the user an intuitive and useful way of quickly separating objects in areas of locally high inter-object occlusion.

4.1 Object Displacement

Objects falling within the current balloon volume are simply displaced to its surface along the radius passing from the balloon center through the centerpoint of the object. The object is drawn normally in its new position. In addition, a ghosted trace of the displaced object and its displacement vector is rendered, giving the user a visual cue of the effects of the space distortion on the environment and its objects.

To further aid the user in maintaining a cognitive map of the actual layout of the current scene, the balloon itself is visually represented by a wireframe 3D sphere, the line segments representing its stacks and slices making it possible to

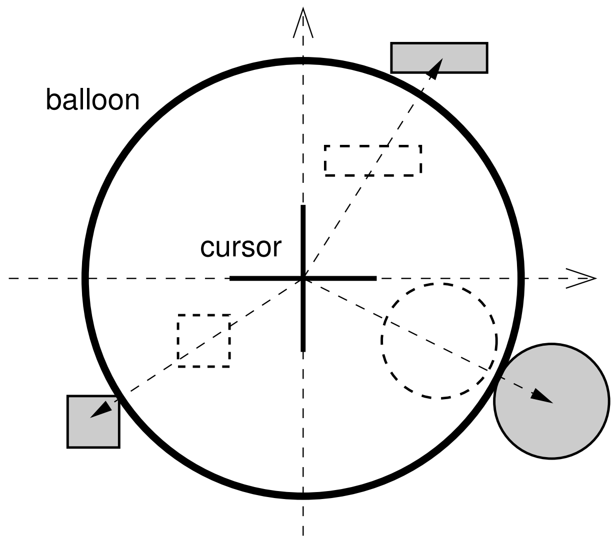

Figure 2: 2D overview of the BalloonProbe technique.

more easily derive the spatial configuration of the displaced objects.

An optional extension to this model is to subject the displaced objects to collision detection so that they do not overlap on the surface of the balloon. The objects are then arranged on the spherical surface using weights proportional to their total displacement distance, conceivably using simple spring equations to derive equilibrium for all objects colliding on the sphere. The result is essentially a force-directed spherical layout. However, this is really only necessary for objects contained by other objects (and, in fact, sharing the same center point), otherwise the user can modify the position and size of the BalloonProbe to disambiguate between these.

4.2 Selective Displacement

In some cases, it may be useful to be able to *selectively displace* some objects while leaving others in their original position, for instance when the user is looking for a particular type of graphical entity in an information visualization application. This is easily done by configuring the BalloonProbe to not affect a certain type of object, identified by some pattern or identifier. Our prototype implementation (see below for more details) can be configured to selectively displace objects based on their color (either making them fixed and the other objects displaced, or vice versa), but using more complex selection criteria is naturally possible. Figure 4 shows an example of this, where all objects except the blue ones have been displaced.

4.3 Implementation

We have implemented a prototype implementation of the BalloonProbe interaction technique on a standard Windows-based PC with 3D graphics hardware. The application is written in C# and uses OpenGL for 3D rendering. Mouse and keyboard is used for input, although we have also experimented with other input devices. It is clear that an interaction technique like this would benefit the most from

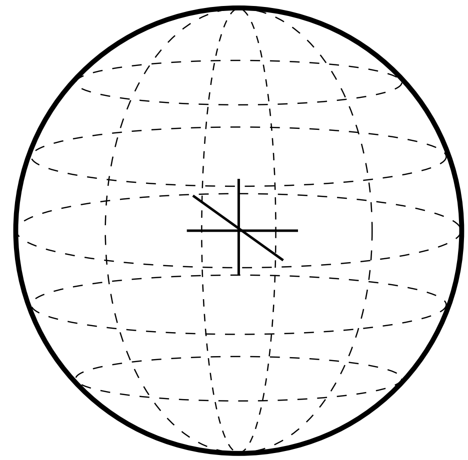

Figure 3: Wireframe representation of the BalloonProbe.

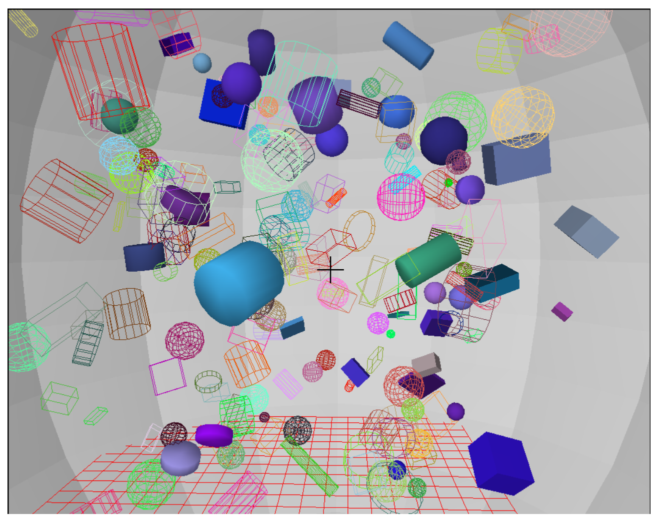

Figure 4: Selective displacement of all non-blue objects.

using a high-end 6DOF motion tracker, like an Ascension Flock of Birds system. Our implementation makes use of the Essential Reality P5 Glove instead, which is a low-budget but extremely cost-efficient 3D tracking solution, resulting in slightly worse accuracy than the high-end alternatives.

Even if the display is not immersive, the application renders a simple 3D cursor, allowing the user to easily control the probe using the input device. The balloon is inflated by pushing and holding a specific button, and deflated when the button is released. Two additional buttons (or a potentiometer) are used to control the size of the balloon. In our implementation, the default inflate toggle is the keyboard control key, and the mouse wheel is used for changing the balloon size.

The test visualization implemented in our prototype is a simple $n \times n \times n$ unit cube populated with various random shapes of configurable density. In addition to the BalloonProbe interactions, the application also supports a simple camera orbit operation around the probe centerpoint to allow the user to easily inspect the displaced objects on the surface of the balloon. Figure 1 shows an example of this 3D environment with a fully inflated BalloonProbe being controlled by the user to study the graphical objects in an area of high congestion.

4.4 User Study

We have conducted an informal user study of the BalloonProbe interaction technique using two volunteers from our research group. The test subjects were asked to experiment independently with the prototype implementation after a short introduction, and were then tasked with solving a few simple visual search tasks. Their comments and feedback were used to improve the prototype. In general, the test subjects indicated that the BalloonProbe was an intuitive and useful tool for quick and effortless occlusion reduction with straightforward and natural interaction.

In the future, we anticipate conducting a formal user study of the BalloonProbe technique using an environment of simple 3D objects (similar to Figure 1) where test subjects are asked to perform visual search of specific objects in the environment. By comparing correctness and completion times with and without the use of the BalloonProbe, we hope to show that performance is improved using our new technique.

5. DISCUSSION

The BalloonProbe technique addresses occlusion reduction by distorting the display space and wrapping objects onto the surface of a spherical force field. Of course, this is simply a projection from a 3D to a 2D space, essentially losing one dimension for the visualization. However, the key feature of the method is that the spherical balloon surface is optimized for occlusion-free inspection through simple camera orbiting. The balloon can either be inflated to a size where no overlap occurs, or force-directed collision detection can be imposed on the displaced objects to avoid overlap altogether.

By displacing objects from their original position in the 3D environment, we may lose vital information encoded in the relative positioning of objects in a visualization. This is naturally a weakness of the BalloonProbe, but we hope to allay this through the rendering of ghosted objects and their displacement vectors. Furthermore, the user is encouraged to alternatively inflate and deflate the balloon multiple times when studying a visualization to note how the individual objects are smoothly displaced.

The BalloonProbe is of special interest in an immersive virtual environment, where the user can use an egocentric viewpoint to actually position herself in the center of the balloon and study it from the inside, getting an excellent view of the displaced objects as well as any objects who were not affected by the selective displacement mechanism of the technique. Thus, in such environments, the probe could be attached to either the body or the manipulator of the user, whereas for non-immersive platforms an exocentric view will work best.

So far, we have described the BalloonProbe as being an invisible or transparent (wireframe) spherical force field that repels objects in the 3D environment. However, for some situations it might be beneficial to render the balloon as solid geometry to further help users separate distracting objects from important ones (see Figure 4 for an example).

6. CONCLUSIONS

We have presented an intuitive space distortion technique for interactively reducing inter-object occlusion in 3D environments. The technique, called BalloonProbe, uses a spherical force field that can be smoothly inflated and deflated just like a balloon, and which will displace any objects it collides with as it expands. By connecting this probe to a 3D-tracked cursor, users can easily pinpoint areas of high congestion and inflate the balloon in this position to separate occluding objects from each other. Ghosted outlines of the objects remain in their original position, allowing the user to maintain a cognitive map of the environment when using the technique. Additionally, users can configure the probe to selectively displace only certain object types while not affecting others, making it possible to quickly clear away irrelevant objects from important objects in a visualization application.

Future work on the BalloonProbe involves testing the interaction technique in a fully immersive virtual environment, as well as trying to integrate it with new or existing VR applications. It would be especially useful to be able to provide an implementation of the BalloonProbe as a standard interaction component for rapid deployment in such applications.

7. REFERENCES

[1] T. T. Elvins, D. R. Nadeau, and D. Kirsh. Worldlets – 3D thumbnails for wayfinding in virtual environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology 1997*, pages 21–30, 1997.

[2] S. Fukatsu, Y. Kitamura, T. Masaki, and F. Kishino. Intuitive control of “bird’s eye” overview images for navigation in an enormous virtual environment. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology 1998*, pages 67–76, 1998.

[3] G. W. Furnas. Generalized fisheye views. In *Proceedings of the ACM CHI’86 Conference on Human Factors in Computer Systems*, pages 16–23, 1986.

[4] J. Lamping and R. Rao. The Hyperbolic Browser: A focus + context technique for visualizing large hierarchies. *Journal of Visual Languages and Computing*, 7(1):33–35, 1996.

[5] J. D. Mackinlay, G. G. Robertson, and S. K. Card. The Perspective Wall: Detail and context smoothly integrated. In *Proceedings of the ACM CHI’91 Conference on Human Factors in Computing Systems*, pages 173–179, 1991.

[6] R. Rao and S. K. Card. The Table Lens: Merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. In *Proceedings of the ACM CHI’94 Conference on Human Factors in Computing Systems*, pages 318–322, 1994.

[7] R. Stoakley, M. J. Conway, and P. Pausch. Virtual Reality on a WIM: Interactive worlds in miniature. In *Proceedings of the ACM CHI’95 Conference on Human Factors in Computing Systems*, pages 265–272, 1995.

[8] N. Wong, M. S. T. Carpendale, and S. Greenberg. EdgeLens: An interactive method for managing edge congestion in graphs. In *Proceedings of the IEEE Symposium on Information Visualization 2003*, pages 51–58, 2003.