COMPUTER GRAPHICS OPTIQUE
Optical Superposition of Projected Computer Graphics

Aditi Majumder and Greg Welch

Department of Computer Science,
University of North Carolina at Chapel Hill, NC, USA
{majumder,welch}@cs.unc.edu

Abstract. We present some ideas and demonstrations for a hybrid projector-based rendering and display technique we call *Computer Graphics Optique*. Instead of partially overlapping projected images to achieve a wide-area display, we completely overlap projected images on top of each other to achieve the addition of light and color in an "optical composition buffer." The idea is to use the optical composition to replace some analytical computation, to increase rendering speed, gain flexibility, intensity range, and intensity resolution. Where projector-based displays are appropriate, potential uses include the optical realization of certain effects normally requiring a digital accumulation buffer, the optical composition of heterogeneous lighting techniques, and the ability to use heterogeneous graphics engines, in parallel. In addition one can make use of the optical projector control of focus augmented with the optical superposition to achieve effects that are otherwise computationally expensive. We believe that this technique offers the possibility of a new paradigm for combined rendering and projector-based display.

1 Introduction

Light projectors have been used in computer graphics for almost as long as graphics has been around. During all of this time they have been used almost exclusively to construct wide-area and high-resolution tiled displays [17, 7, 10, 15, 13]. A typical design goal for such systems is to maximize the usable display area, and so overlap between projectors is minimized. Recently however we wondered what might be gained by *maximizing* the overlap between multiple projectors—that is, by "stacking" or *superimposing* the projector images on top of each other to achieve addition of light and color *optically*.

For example, this optical superposition certainly yields brighter displays. Taking advantage of this property, many popular projector vendors like Sony, Panasonic, Epson and others are now providing features to 'stack' 2-3 projectors to achieve two or three times brighter images. They are also providing lens shift controls to facilitate the alignment of imagery from multiple projectors right on top of each other.

It appears that optical superposition of imagery was being done in a limited fashion even a century ago. In 1888 Emile Reynaud created a machine called *Theatre Optique* (Figure 1), a revised version of his earlier invention, the *praxinoscope*. In Theatre Optique a subsidiary light called "magic lantern," was used to project a stationary background on the same screen where the action film was projected. In more recent times,

Fig. 1. Theatre Optique

the flight simulator community used a similar method called *calligraphic lights* to superimpose high intensity and brighter runway lights on a projected scene.

Beyond increased brightness, we believe optical superposition augmented by some optical effects like blurring (available in most projectors), can be used in modern computer graphics to replace expensive *computation*, to increase rendering *speed*, to achieve *flexibility*, and to increase intensity *range* as well as *resolution*. We like to call such a hybrid projector based rendering and display technique *Computer Graphics Optique* (CGO) in deference to Reynaud's Theater Optique.

At this point, one might reasonably wonder about achieving similar effects with specialized graphics hardware, for example by compositing results from multiple rendering boards like the PixelFlow machine [16], using interconnected parallel graphics cards like Quantum3D [1], or composition/interconnect systems such as the Lightning-2 system at Stanford University [5]. What we hope to convey is at the very least an interesting *alternative* to specialized or monolithic graphics engines when projector-based displays are appropriate, an alternative that leverages the effective "bandwidth" into the "optical buffer," i.e. the projector screen or display surface.

In this paper we present a number of ideas for the uses of optical superposition in projector-based interactive computer graphics. While we have experimented with all of the ideas, some are more well developed than others. For these we have done some preliminary analysis of their usefulness, and we have developed some relatively simple methods for implementing them. The practical utility of some of the less-developed ideas are yet to be proven, however we also present those ideas with the hope of sharing them with the graphics community to obtain feedback that will help us to better assess their utility and target our efforts.

2   Interactive Depth of Field via Optical Blurring

We believe there is a significant advantage to using the optical controls on a pair of projectors to achieve computationally-intensive realistic effects like blurring at *inter-*

active rates, which is not currently possible even using specialized graphics hardware. Using the projector optics one can achieve *depth of field* (DOF) effects at interactive rates, something that would add valuable realism to projector-based environments for flight simulation and training, scientific visualization, immersive virtual environments, architectural walkthroughs, and entertainment. Beyond a lack of realism, the absence of DOF effects leads to eye strain as *fusion stimuli* are generated for objects outside the limited depth of focus of the eye. DOF effects are also sometimes used as depth cues. For example for night-vision training of helicopter pilots, pilots rely on DOF cues near the ground to compensate for the narrow fields of view of night vision goggles. By some simple experiments, we demonstrate how optical blurring and superposition can be used to achieve a DOF effect in real-time.

2.1   Algorithm

We use two projectors whose images overlap completely. We designate one projector as the *focused* projector and the other as the *defocused* projector. We then adjust the optical focus of the latter to correspond to the maximum circle of confusion for the application. At run-time we use a two pass rendering approach. After the image is rendered by the traditional graphics pipeline in the first pass, each pixel is divided into two complementary components based on its distance from the focal plane in the second pass. Thus we get two images: one to be sent to the focused projector, and the other to the defocused projector. When projected, the two overlapping images combine via optical superposition, resulting in a single image with depth-dependent blur (Color Plate).

In effect we use the projector optics to perform two samples of the image convolution, and then we use optical superposition to blend a weighted combination of the results. We can vary the depth of the focal plane in real time. One could do so based on a user's head position for example, or ideally (if available) eye tracking could be employed [18] to adjust the focal plane.

We can use a similar technique to achieve a focused inset region in a blurred 2D image (or visa versa). We use a focused projector to render the inset region, and a defocused projector to render the rest. We can move the focused inset around in the image interactively by moving a pattern in the alpha buffers (Color Plate).

2.2   Analysis

In this section, we analyze what we are actually getting using this method and how similar it is to the ideal depth of field effect.

Depth of Field Effect : Based on a *thin lens model* for the human eye, Figure 2 helps explain the blurring of the retinal image of points beyond the limited DOF. We let r be the distance between the lens and the retina, e be the aperture of the eye lens, and b be the focal or fixation depth. An object point at depth b will form a focused spot (circle of radius 0) on the retina, while an object point P at some other depth z, will form a (blur) circle of radius c on the retina. In general the *thin lens equation* relates focal length, the distance to the object point, and the distance to the focused image as

Fig. 2. The Depth of Field Effect in Human Eye

$$\frac{1}{object\_dist} + \frac{1}{image\_dist} = \frac{1}{focal\_length} \tag{1}$$

Using this expression, the principle of *similar triangles*, and the fact that an object at depth z focuses at a distance d from the lens, we can obtain the following expression for the radius c of the corresponding circle of confusion in Figure 2:

$$c_z = |er(\frac{1}{b} - \frac{1}{z})|. \tag{2}$$

Relative Radius of Confusion : Let us assume we are rendering a scene with near and far plane at depth n and f respectively. We can say that the *maximum* radius of a circle of confusion, c_m, will occur for an object point on the far plane when the focal plane is coincident with the near plane, or vice-versa. Hence, from Equation 2

$$c_m = \left| er \left(\frac{1}{n} - \frac{1}{f} \right) \right| \tag{3}$$

We define the *relative radius of confusion*, c_r, for a point at depth z as the ratio of c_z and c_m as follows.

$$c_r = \frac{c_z}{c_m} = \left| \frac{\frac{1}{b} - \frac{1}{z}}{\frac{1}{n} - \frac{1}{f}} \right| \tag{4}$$

Notice that c_r is zero for an object point at b and increases as the distance of the object from the fixation depth increases.

Second Rendering Pass : In the traditional graphics pipeline, the perspective transform z' of depth values z for a 3D scene with a near and far plane at $z = n$ and $z = f$ respectively, is given by

$$z' = \frac{f(z - n)}{z(f - n)} \tag{5}$$

Let us assume that we are rendering a scene when the focal depth is $z = b$. It turns out that relative radius of confusion c_r for any z can be obtained by simply subtracting b',

the perspective transform of b, from z' obtained from the z-buffer.

$$|z' - b'| = |\frac{f(z-n)}{z(f-n)} - \frac{f(b-n)}{b(f-n)}| = \frac{|\frac{1}{b} - \frac{1}{z}|}{|\frac{1}{n} - \frac{1}{f}|} = c_r \qquad (6)$$

During the first pass at runtime we copy the *desired image* (just rendered) in texture memory. In the second pass, we read the z-buffer values and then generate an *alpha mask* for the defocused projector by calculating the c_r at every pixel by subtracting b' from the values read from the z-buffer. We simultaneously generate a complementary alpha mask for the focused projector. Finally we use *alpha blending* and texture map the desired image on the projector's image plane using the respective alpha masks.

Summary : Thus, an ideal DOF effect is generated in a graphics systems by blurring each pixel by their relative radius of confusion. However, our method achieves this effect by the constant radius of confusion produced by the defocused projector while limiting the amount of light from the defocused projector to be proportional to the relative radius of confusion. We project the rest of the light from the focused projector with the hope that this will perceptually create a circle of confusion equivalent to the ideal one.

3 Parallelism and Flexibility in Rendering

Wile conventional graphics hardware has achieved significant performance in the years, it has done so while effectively restricting interactive computer graphics to specialized APIs and rendering paradigms. It is currently difficult, if not impossible, to match and mix between different rendering styles while rendering the same scene. We believe that optical superposition has the potential to liberate us from this rigidity and give us a new flexibility to simultaneously combine rendering methods for the same imagery. Even for one or more similar rendering methods the technique could potentially be used to exploit parallelism, improving rendering speed by sharing tasks between independent "optical pipelines."

Separating Lighting Components : One universal aspect of computer graphics where superposition inherently plays a fundamental role is virtual scene illumination. Such illumination is typically modeled as a sum of independent terms that are added together. Separating these components offers the flexibility of using different lighting effects in parallel for the same scene. For example, we can optically combine diffuse lighting from a global illumination algorithm with specular lighting from a simple Phong model. As a proof of concept, we compute the diffuse and specular lighting using two different graphics pipeline, and project the images asynchronously on top of each other from two different projectors. One might extend this to separate lighting into different spatial regions and resolutions also. For example, lower resolution computation and projectors for ambient or diffuse lighting, and higher resolution for specular effects or hard shadows.

Stages in Graphics Pipeline : There is often no reason for texture mapping to be coupled to lighting, which normally depends only on the surface normal and material properties. In some cases each could be computed with a different platform or pipeline, and projected asynchronously.

Multi-Pass Rendering : Many algorithms have multiple rendering passes which are independent of each other, and the *accumulation buffer* [9] is used to add colors. All such algorithms can have their different rendering passes implemented in multiple graphics pipeline in parallel, each being output to a projector, and the results can be combined in an *optical accumulation buffer*. For example, Diefenbach achieves effects like shadows, reflections and refractions in [4] using an accumulation buffer and a multi-pass rendering method. Or, for example, sometimes software and computational resources place practical limits on the number of lights that can be used in a scene (OpenGL is limited to eight lights). Thus an interactive ray-traced scene with 50 lights would probably take multiple rendering passes. Here also, we can render subsets of lights in parallel on separate graphics pipelines, and then optically combine the (projected) results (Color Plate) to achieve highly specular and photorealistic scenes without compromising on the interactive rates.

Synchronization Issues : In Bastos's *superposition rendering* [2], where all lighting computation is sequential and synchronized, Bastos notes that the computation of the glossy reflections is the primary rendering bottleneck. While we believe we can help by decoupling tasks, rendering them in parallel, and re-combining them optically, we need to develop an understanding of the perceptual effect of such asynchronous rendering. We are unable to find directly relevant user studies in the literature, and therefore intend to undertake this investigation.

4 Higher-Fidelity Imagery

The color range in conventional graphics ranges only from $0 - 255$. Here we show that a display made up of superimposed images from multiple projectors is not only brighter but has a higher intensity resolution which can be used to render outdoor/natural scenes more realistically.

A display d capable of producing a lowest intensity $I(0) = L$ (the black level) and highest intensity $I(n - 1) = H$ (the white level), is characterized by the following.

- *Range:* The difference between the highest and the lowest intensity levels, $H - L$.
- *Dynamic Range:* The ratio of the highest and lowest intensities, $\frac{H}{L}$.
- The *number of levels*, n, between the highest and lowest intensities, which is typically 256. This is the number of unique discrete colors that can be produced.
- *Intensity Resolution:* Levels per unit intensity, $\frac{n}{H-L}$. This is true only for devices with a linear response. Thus, we assume that all our projectors have been corrected for different kind of non-linearities [10].

With the use of multiple overlapping projectors, the *range* is the sum of the individual projector ranges. The *dynamic range* is approximately the same as that of the

individual projectors. It can be shown that a CGO display made up of m projectors each having n discrete intensity levels can have a minimum of $O(mn)$ and a maximum of $O(n^m)$ discrete intensity levels. As a result, we get a higher *intensity resolution* from a CGO display.

The number of levels (n) that is sufficient for a *monochrome* display is given by $1.01 = (H/L)^{1/n}$ [8]. Based on the response curves of multiple projectors [11] we find that $n = 500$ would be sufficient for a monochrome projector. Further, [3] says that humans are $3 - 4$ times more sensitive to luminance differences in the presence of color. Thus, the problems arising from the low value of n offered by current projectors can be addressed to some extent by CGO. Further, the larger range will result in higher quality images. We demonstrate this by using three overlapped projectors to render a high dynamic range scene (Color Plate).

5 Limitations

The CGO idea is not without limitations. The most significant limitation is that it is only useful when a projector-based display is appropriate. Further, since today's projectors do not produce zero intensity for a corresponding zero input, optical superposition increases perceived "black level." However, it is reasonable to expect this to improve with new technology [6]. A potential limitation when using CGO with parallel optical pipelines is that one might need to explicitly manage certain occlusions, instead of relying on the hardware to resolve them. Finally, while there may be concerns about luminance/color uniformity and geometric blending issues, we have used the methods of [11] to achieve results that are comparable to a single-projector setup.
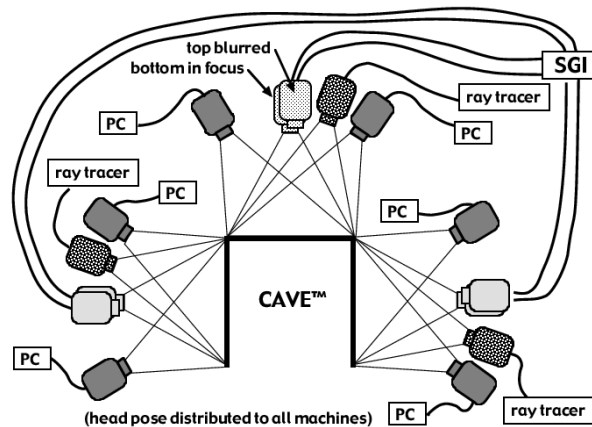

**Fig. 3.** A hypothetical CGO CAVE setup.

6 Conclusion

To summarize, we believe that optical superposition (CGO) offers a new hybrid rendering and display paradigm for immersive virtual environments (Figure 2), one that replaces some computation, offers greater rendering flexibility, and facilitates rendering parallelism. In support of this belief we have conducted some simple demonstrations of a few CGO applications. It appears at the least that the most valuable use is the generation of DOF effects at interactive rates for projector based systems. However we do also believe that the remaining CGO applications we have identified (and others we have not) can be used to achieve higher-fidelity projector-based imagery. We anticipate finding and/or developing a mathematical framework appropriate for the additive superposition of light, so that we can more rigorously find and characterize the uses and usefulness of CGO.

References

1. "Mercury: Realtime anti-aliasing 3d graphics sub-system," Technical Report, http://www.quantum3D.com/products%20pages/mercury.html, [cited Jan 2000].
2. Rui Bastos. *Superposition Rendering : Increased Realism for Interactive Walkthroughs*. PhD thesis, University of North Carolina at Chapel Hill, Department of Computer Science, 1999.
3. R.A. Chorley and J. Laylock, "Human factor consideration for the interface between electro-optical display and the human visual system," In *Displays*, volume 4, 1981.
4. Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A.Defanti, "Surround-screen projection-based virtual reality : The design and implementation of the cave," In *Proceedings of ACM Siggraph*, 1993.
5. "Flash : Graphics System and Architecture,"Department of Computer Science, Stanford University, [cited on Jan 2001] Available from http://graphics.stanford.edu/projects/flashg/,
6. White Pages of ColorVision, http://www.colorvision-lasers.com/, [cited Jan 2001].
7. P. Diefenbach and N. Badler, "Pipeline rendering : Interactive refractions, reflections and shadows," In *Displays(Special Issue on Interactive Computer Graphics)*, 1994.
8. J.D. Foley, A. Van Dam, S.K. Feiner, and J.F.Hughes, *Computer Graphics Principles and Practice*. Addison Wesley, 1990.
9. Paul E. Haeberli and Kurt Akeley, "The accumulation buffer: Hardware support for high-quality rendering," *Computer Graphics (Proceedings of SIGGRAPH 90)*, 24(4):309–318, August 1990.
10. G. Humphreys and P. Hanrahan, "A distributed graphics system for large tiled displays," In *Proceedings of IEEE Visualization*, 1999.
11. Aditi Majumder, Zue He, Herman Towles, and Greg Welch. "Achieving Color Uniformity Across Multi-Projector Displays," In *Proceedings of Visualization*, 2000.
12. Thomas L. Martzall, "Simultaneous raster and calligraphic crt projection system for flight simulation," In *SPIE Proceedings,Electroluminescent Materials, Devices, and Large-Screen Displays*, volume 1910, 01/31/1993 - 02/05/1993.
13. R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs, "The office of the future: A unified approach to image based modeling and spatially immersive display," In *Proceedings of ACM Siggraph*, pages 168–176, 1998.
14. Ramesh Raskar, "Immersive planar displays using roughly aligned projectors." In *Proceedings of IEEE Virtual Reality 2000*, 1999.
15. Samanta Rudro, Jiannan Zheng, Thomas Funkhouse, Kai Li, and Jaswinder Pal Singh, "Load balancing for multi-projector rendering systems," In *SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August 1999.

16. S.Molnar, J. Eyles, and J. Poulton, "Pixelflow : High-speed rendering using image composition," In *Proceedings of SIGGPARH*, pages 231–240, 1992.

17. M. Lacroix, "A HDTV Projector for Wide Field of View Flight Simulators," Presented at *IMAGE VI Conference, Scottsdale, AZ*, July 14–17, 1992.

18. K. Talmi, J. Liu, "Eye and Gaze Tracking for Visually Controlled Interactive Stereoscopic Displays," Proceedings of *Image Communication*, 1998.

Computer Graphics Optique : Color Plate

a. Two superimposed projectors. Both are in focus.

b. One of the projectors are defocussed. The green teapot at the back is in focus.

c. This is the contribution of the defocussed projectors in b.

d. This is the contribution of the focussed projectors in b.

e. The silver teapot in the middle is in focus now.

Highly Specular Scene : Top - One projector projecting the effects of 7 highly specular lights. Middle - Addition of one more projector projecting 7 different specular light sources. Bottom - A third projector projecting 6 more light sources.

The circular inset is projected from one projector and the rest from another. Left : Both the projectors are in focus . Right : One of the projectors is defocussed.

Left : A high dynamic range scene (1000:1) being projected by one projector. Right : The same scene being projected by a CGO display made up of three projectors. *(The pictures do not show the differences very well because of the inherent limitations of the involved devices like the camera and the printer. These should have similar color range and resolution to reproduce the result.)*