AIM

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

Centimeter

Inches

MANUFACTURED TO AIM STANDARDS
BY APPLIED IMAGE, INC.

Title:

Fast Stereoscopic Images with Ray-Traced
Volume Rendering

Author(s):

Charles D. Hansen
Stephen J. Adelson

Submitted to:

ACM Workshop on Volume Visualization

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy.

MASTERForm No. 836 R5
ST 2629 10/91

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Fast Stereoscopic Images with Ray-Traced Volume Rendering

Stephen J. Adelson

Los Alamos National Laboratory

MS M986

Los Alamos, New Mexico 87545

e-mail: adelson@lanl.gov

Charles D. Hansen

Los Alamos National Laboratory

MS B287

Los Alamos, New Mexico 87545

e-mail: hansen@acl.lanl.gov

ABSTRACT

One of the drawbacks of standard volume rendering techniques is that it is often difficult to comprehend the three-dimensional structure of the volume from a single frame; this is especially true in cases where there is no solid surface. Generally, several frames must be generated and viewed sequentially, using motion parallax to relay depth. Another option is to generate a single stereoscopic pair, resulting in clear and unambiguous depth information in both static and moving images.

Methods have been developed which take advantage of the coherence between the two halves of a stereo pair for polygon rendering and ray-tracing, generating the second half of the pair in significantly less time than that required to completely render a single image. This paper reports the results of implementing these techniques with parallel ray-traced volume rendering. In tests with different data types, the time savings is in the range of 70 - 80%.

Fast Stereoscopic Images with Ray-Traced Volume Rendering

Stephen J. Adelson

Los Alamos National Laboratory

MS M986

Los Alamos, New Mexico 87545

e-mail: adelson@lanl.gov

Charles D. Hansen

Los Alamos National Laboratory

MS B287

Los Alamos, New Mexico 87545

e-mail: hansen@acl.lanl.gov

1. INTRODUCTION

Ray-tracing is a common method for generating realistic images of complex surfaces. A similar technique, sometimes referred to as ray-casting, is also frequently used with volumetric data. Rays are cast through the volume from each pixel, accumulating color and opacity as they travel, and returning an overall color for that position. Like all volume rendering techniques, however, a single image can be difficult to interpret because of image artifacts which hinder structure determination. The problem is even more complicated when there is no solid surface and color is used to represent some characteristic of the data rather than a lighting function. Real-world depth cues such as luminosity, perspective, shading, and occlusion are, at best, difficult to use in many volumetric images, particularly those with no isosurface [Adel94b].

The general solution to this problem is to generate multiple images of the same data from displaced view points and view them sequentially. Using motion parallax in this way is an effective depth cue, but it has two major shortcomings. First, multiple images of complex data involve significant allocation of both time and computing resources. Second, depth information is only available while the image is in motion. Ideally, a volume could be paused in its movement and examined. If the progression of views is halted, the motion depth cue is lost.

Another solution to the same problem is to generate a single stereoscopic pair. Stereoscopy gives unambiguous depth information when it is the only available depth cue, and depth discrimination is enhanced when other cues are also included [ReBS90, YeSi90, HoMc91,

Rein91, SoMi91, AABH93, WaAB93, HBCP94]. Only two images need to be rendered, the two halves of the stereo pair, and no motion is necessary; static images can be examined with no loss of perceived structure. In this paper, it is shown that both halves of a stereoscopic pair of ray-traced volumetric images need not be fully rendered. Exploiting the coherence between the two views, the second half of the stereo pair can be generated in a fraction of the time of the first half.

2. BACKGROUND AND PREVIOUS WORK

2.1 Stereoscopic Reprojection

Most volume renderers use parallel projections for generating images. In this section, stereoscopic projection is derived for parallel projections, although it is a fairly straightforward matter to derive for perspective projection [HoMc93, Adel94a].

The most efficient viewing geometry for parallel stereoscopic imaging is illustrated below in Figure 1. There are two view points, separated by a distance e : the left center of projection (LCoP) for the left-eye image and the right center of projection (RCoP) for the right-eye image. Instead of placing the centers of projection symmetrically about the origin (as they would be in traditional rendering geometries), one is positioned at the origin and the other at the point $(e \cos(\Phi/2), 0, e \sin(\Phi/2))$. This is accomplished by changing the viewing transformation matrix so that the center of the viewing coordinate axes is at $(e/2 \cos(\Phi/2), 0, e/2 \sin(\Phi/2))$ and the w-axis of the (u, v, w) viewing coordinate system is $(-\sin(\Phi/2), 0, \cos(\Phi/2))$. This makes the center of rotation $P = (0, 0, R)$, where $R = e/[2 \sin(\Phi/2)]$. No additional calculations are necessary to view-transform data when this change is made in the viewing transformation matrix.

The left-eye projection is simply a parallel projection based on the data position after the viewing transformation. The right-eye projection involves a rotation of $-\Phi/2$ about the y axis, a translation of $(-e, 0, 0)$ to place the RCoP at the origin, and a second rotation of $-\Phi/2$ about the y axis to place P back on the z axis. Using c and s as abbreviations for $\cos(\Phi/2)$ and $\sin(\Phi/2)$ respectively, the projection matrix for the right-eye view is:

$$\begin{array}{ccc}
\begin{array}{ccc} c & 0 & s & 0 \\ 0 & 1 & 0 & 0 \\ -s & 0 & c & 0 \\ 0 & 0 & 0 & 1 \end{array} & \begin{array}{ccc} 1 & 0 & 0 & -e \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{array} & \begin{array}{ccc} c & 0 & s & 0 \\ 0 & 1 & 0 & 0 \\ -s & 0 & c & 0 \\ 0 & 0 & 0 & 1 \end{array} \\
\text{rotate } -\Phi/2 & \text{translate } -e & \text{rotate } -\Phi/2
\end{array} = \begin{array}{ccc} c^2-s^2 & 0 & 2cs & -(c \ e) \\ 0 & 1 & 0 & 0 \\ -2cs & 0 & c^2-s^2 & (s \ e) \\ 0 & 0 & 0 & 1 \end{array}$$

projection matrix

Figure 1: Displaced Parallel Projection Geometry

So the view-transformed point (x_p, y_p, z_p) projects to (x_{sl}, y_{sl}) and (x_{sr}, y_{sr}) , where

$$x_{sl} = x_p \quad (1)$$

$$\begin{aligned}
x_{sr} &= x_p [\cos^2(\Phi/2) - \sin^2(\Phi/2)] + z_p [2 \cos(\Phi/2) \sin(\Phi/2)] - e \cos(\Phi/2) \\
&= x_p \cos(\Phi) + z_p \sin(\Phi) - e \cos(\Phi/2)
\end{aligned} \quad (2)$$

$$y_{sl}, y_{sr} = y_p \quad (3)$$

Since $e \cos(\Phi/2)$ is a constant term, both projections can be calculated using two additions and two multiplications.

2.1.2 Limitations on Viewing Angle or Volume Scale in Parallel Projections

In parallel projected stereoscopic images, horizontal parallax $(x_{sr} - x_{sl})$, the distance between projections) is unbounded. Horizontal parallax should certainly be limited to the interocular distance so that one need not become walled in order to view the images. Furthermore, while experienced viewers can fuse the interocular parallax [Hodg93], most casual viewers can tolerate much less [Hodg92]. To keep horizontal parallax in a reasonable range, the z

values of the data and values of the rendering constants e and Φ must be kept within a specific range.

Let $e_{\min}$ and $e_{\max}$ be the near (negative parallax) and far (positive parallax) limits of horizontal parallax. These can be described as a function of the physical distance from the viewer to the screen, D , expressed in screen units:

$$e_{\min} = -K1 D \quad (4)$$

$$e_{\max} = K2 D \quad (5)$$

$K1$ and $K2$ are positive constants which are generally 0.028 or less [Hodg92]. Given that a separation of $e_{\min}$ should occur at the closest z point of the data and $e_{\max}$ at the farthest, a range for z can be derived [Adel94a]:

$$D (1 - K1/[2 \sin(\Phi/2)]) \leq z \leq D (1 + K2/[2 \sin(\Phi/2)]). \quad (6)$$

If D is fixed at some comfortable viewing distance, the range of z may be matched to the data by manipulating the viewing angle Φ . Alternately, the data may be rescaled to accommodate any desired viewing angle. The value of e can be found by recognizing that the distance to point P , at which there is no parallax, is equal to the distance to the screen, D , so as can be seen in Figure 1:

$$e = 2D \tan(\Phi/2) \quad (7)$$

2.1.3 Parallel Shears

Since Φ is usually small, it has been suggested that the substitutions $\sin(\Phi) = \Phi$ and $\cos(\Phi) = 1$ could be used with little loss of accuracy [Lips79]. Even at 6.5 degrees, more than four times the larger recommendations for e , the error is less than 0.65%. Using these simplifications reduces the x coordinate calculations as follows:

$$x_{sl} = x_p \quad (8)$$

$$x_{sr} = x_p + z_p \Phi - e \quad (9)$$

Both pairs of coordinates can be found in two additions and one multiplication using displaced parallel projection. The same considerations about unbounded parallax exist for shears as they do for rotations.

2.2 Ray-Traced Volume Rendering

In ray-traced volume rendering, such as that described by Levoy [Levo88], rays are traced through the volume, accumulating values and opacities at intervals (regular or irregular) until an opacity limit is reached (often 1.0) or the ray exits the volume. The colors and opacities are found by trilinear interpolation from the eight nearest voxels.

Ray samples are accumulated by the use of Porter and Duff's **over** operation [PoDu84]:

$$S(i) \text{ over } S(j) = S(i) + (1 - \alpha(i)) S(j) \quad (10)$$

where $S(i)$ is a color / opacity pair $[C(i), \alpha(i)]$ representing a portion of the ray closer to the viewpoint than the second pair $S(j)$.

Note that when rays are sampled in regularly spaced intervals, projection from the left-eye to the right becomes even simpler. Looking at equation (9), the projection value on a ray will increase at each interval by $\Delta z \Phi$ (or $\Delta z \sin(\Phi)$, if rotations are used instead of shears), which is a predetermined constant. Furthermore, if the left-eye rays are processed sequentially along scan-lines, the initial right-eye projection from two consecutive left-eye rays will differ by the constant Δx (or $\Delta x \cos(\Phi)$). On average, then, the projection value can be found with a single addition operation.

2.2.1 Parallel Volume Rendering

Parallel architectures are sometimes used for volume rendering because of their speed and the memory for rendering huge data sets which is not available on traditional architectures. Ma, et al. described an algorithm for parallel volume rendering using a data distributed model on the Thinking Machines CM-5 [ThMC91, MPMK93]. In their implementation, the volume is divided among a power of two number of nodes, using a k-D tree structure with binary subdivision along successive world axes. Each subvolume is transformed to correspond to the viewing specifications and given a bounding box orthogonal to the world axes, through which rays are cast.

Rays accumulate values at regular intervals, which are dependent on volume size and not image size. While load-balancing problems remain to be addressed, the method has good characteristics for accelerated rendering. The algorithm presented here uses Ma, et al.'s code as a base for the renderer. The displaced parallel projection with shears, as described in

section 2.1.3, is used. It should be noted, however, that the stereoscopic algorithm does not require a parallel renderer, and would work just as well on a serial computer.

2.3 Guaranteeing Correct Stereoscopic Views

If rays are allowed to traverse completely through the volume, the samples computed for the left eye are projected to the right-eye viewpoint, resulting in a correct stereoscopic view. However, a common method for accelerating the volume rendering process is early ray termination. When the accumulated ray opacity reaches 1.0, the remaining samples encountered by the ray contribute nothing to the final color.¹ Thus, it is possible to terminate the ray traversal at this point without any loss of information. Yet if the left-eye ray is terminated early, samples are ignored which might contribute to one or more right-eye rays. It is necessary to handle this case properly if a correct stereoscopic view is to be obtained.

Figure 2 shows parallel rays from the two view points with the rotational angle greatly exaggerated for clarity. From this figure it can be seen that if rays in the left-eye view are evaluated right to left along scan-lines and front to back along the ray, the right-eye rays will accumulate values in a front to back order as well. As the left-eye rays accumulate values, the right-eye projection of the current sample is calculated, and if the appropriate right-eye ray has not reached full opacity, the sample is accumulated into that ray as well.

Suppose, as in Figure 2, that a left-eye ray (ray 3) accumulates full opacity and terminates at the round solid object. Nothing is known about the data which may lie beyond this termination point. It is possible that samples which accumulate to a left-eye ray processed after the terminated ray (rays 4 - 6) would project into a right-eye ray whose view may be obscured by the unknown area (rays r2 - r4). To avoid this, the maximum allowable right-eye projection will be the final right-eye projection of any left-eye ray already evaluated. In this example, as ray 4 is processed, no right-eye projection will be allowed to the right of ray r1, since there may be unknown material beyond the termination of ray 3 which was never accumulated into the right-eye rays.

¹If isosurfaces exist in the volume rendering, the opacity reaches 1.0 when a surface is encountered. Early ray termination can be thought of as the existence of occluding objects at the voxel where the opacity reaches 1.0.

Figure 2: The round object terminates ray 3 and projects to ray r1 in the right eye. The oval object in ray 4 would project to ray r2, but that projection will not be allowed since an object (like the square) may exist past the termination point in ray 3. No right-eye projection is allowed farther to the right of the final projection of any left-eye ray already evaluated.

When all left-eye rays have been calculated, the right-eye rays are re-examined. Any right-eye ray which has not accumulated full opacity and whose last projection is not on the far side of the volume bounding box, indicates some left eye ray was terminated early. In order to generate the correct stereoscopic view, it is necessary to continue the ray traversal from the last projection into the ray. This algorithm assures that right-eye rays accumulate only those samples from the left-eye view which are guaranteed to be visible and that the right-eye view evaluates voxels only when and where it must.

3. PERFORMANCE TESTS AND DISCUSSION

A modified version of Ma, et al.'s parallel volume renderer was created which would project voxel values using the displaced parallel projection. Three sets of volumetric data were used to test the renderer: a CT scan of a human head (128^3 - Figure 3), a Magnetic Resonance Angiography (MRA) data set showing the vascular structure in the brain of a patient ($256 \times$

256 X 128 - Figure 4), and a 256^3 CFD data set computed on the CM-200 showing the onset of turbulence (Figure 5). Each set was evaluated at three different rendering resolutions (single image sizes of 64 X 64, 128 X 128, and 256 X 256) and five different node configurations (32, 64, 128, 256, and 512 nodes).

Figure 3: Stereo pair of the head data. The right eye image was generated using the technique described in this paper.

Figure 4: Stereo pair of the brain vessel data. The right eye image was generated using the technique described in this paper.

Tables 4, 5, and 6 below display the savings of the algorithm in three ways. First, the actual time savings, comparing the time to generate a full right-eye image with the fast stereoscopic algorithm time to generate the second image. This savings will be measured on a scale of 0 to 100%. The second measure is the percentage of left-eye voxels retained in the right eye view. The measure counts only those voxels which actually accumulate into right-eye rays; voxels which project to a terminated ray are ignored. The final measure is the ratio of the

number of voxels evaluated solely for the right-eye to those evaluated for the left-eye (projected to the right eye or not).

Figure 5: Stereo pair of the vorticity data. The right eye image was generated using the technique described in this paper.

Tables 1-3 show the raw numbers of these three measures for the head data, vessel data, and vorticity data respectively. All times are reported in milliseconds, all voxel counts in thousands. Tables 4-6 show the same results as percentages. The percentages are calculated from the raw data as follows: the time savings percentage is $(1.0 - \text{projection time} / \text{full time})$, the retained voxels percentage is $(\text{saved voxels} / \text{left voxels})$, and the voxel ratio percentage is $(\text{right voxels} / \text{left voxels})$

The timing mechanism used for the stereoscopic algorithm is only accurate to the nearest millisecond. As seen in Tables 4-6, the differences in time savings for a particular data set are insignificant for a fixed image size but with different node configurations.² The savings does increase with image resolution. As image size grows, the parallax between projected points remains the same (as it is based on viewer distance, not image size), even though the horizontal size of the image has increased. In effect, this is equivalent to keeping the image size constant and decreasing the projection distances. Since areas of terminated rays tend to appear in clusters, a smaller effective projection distance means that it is less likely that a sample from the left eye will project into a terminated ray-cluster in the right eye. More left-eye samples will then be retained in the right-eye view, and as a result, fewer samples will need to be evaluated solely for the right-eye rays. (Both phenomenon are seen in the results

²Although the rendering time decreases as the node configuration increases, the *savings* of the algorithm remains fairly constant.

as image size increases.) Since the projection of left-eye samples to the right eye is a much faster operation than evaluating voxels for the right-eye, the time savings also increases with image size.

Table 1: Head Data Results - Raw Data (time in msec. and voxels in thousands)

Size	Nodes				
	32	64	128	256	512
64²					
projection time	71	35	20	11	5
full time	277	137	74	39	19
left voxels	203	209	217	217	222
saved voxels	195	202	211	211	217
right voxels	10	9	9	9	9
128²					
projection time	234	114	65	33	15
full time	1084	533	291	153	73
left voxels	813	835	867	867	888
saved voxels	796	822	856	856	878
right voxels	21	18	16	16	16
256²					
projection time	833	413	230	116	54
full time	4285	2104	1146	599	287
left voxels	3249	3340	3466	3466	3551
saved voxels	3215	3312	3443	3443	3531
right voxels	43	36	31	31	29

The number of voxels retained, those samples evaluated for the left-eye view that were reused in the right-eye view, is high in all cases, better than 96%. This value will be larger if the voxels are less opaque, as this will result in fewer terminated rays in the right-eye view. The head data, which contained an isosurface, had the smallest retained voxels percentage since the rays terminated when they reached the surface. In contrast, the vessel data had at least 98.2% retention, and the vorticity data, at least 99.3%. It has already been explained that this percentage will rise with image size; it will also (generally) rise with node configuration size. When the number of nodes increases, the area of the volume which a given node is rendering will decrease. A smaller volume area lowers the probability that a ray passing through the area will terminate, so that the retained voxels percentage increases. However, there is a possibility that poor partitioning of the data will result in several dense blocks, slightly lowering the overall savings. This is seen in the 64 X 64 vorticity data, in going from 64 to 128 nodes (Table 6).

Table 2: Vessel Data Results - Raw Data (time in msec. and voxels in thousands)

Size	Nodes				
	32	64	128	256	512
64²					
projection time	73	38	19	9	4
full time	319	159	83	43	21
left voxels	275	275	277	278	278
saved voxels	270	270	273	275	275
right voxels	8	8	6	7	6
128²					
projection time	260	130	71	34	17
full time	1284	639	333	173	86
left voxels	1125	1125	1132	1137	1137
saved voxels	1113	1113	1125	1132	1132
right voxels	19	19	12	11	11
256²					
projection time	1005	487	259	133	67
full time	5002	2503	1311	667	333
left voxels	4501	4501	4529	4547	4547
saved voxels	4473	4473	4515	4539	4539
right voxels	42	42	25	19	19

Table 3: Vorticity Data Results - Raw Data (time in msec. and voxels in thousands)

Size	Nodes				
	32	64	128	256	512
64²					
projection time	1536	68	38	18	8
full time	6696	321	167	88	42
left voxels	572	573	573	573	574
saved voxels	567	569	570	570	569
right voxels	7	6	7	7	11
128²					
projection time	528	249	133	69	31
full time	2639	1267	657	347	158
left voxels	2288	2292	2294	2294	2296
saved voxels	2275	2285	2289	2289	2290
right voxels	19	12	11	11	13
256²					
projection time	2034	959	500	263	125
full time	10558	4989	2582	1363	653
left voxels	9203	9218	9229	9229	9234
saved voxels	9164	9200	9220	9220	9228
right voxels	55	30	20	20	20

Table 4: Head Data Results - Percentages

Size	Nodes				
	32	64	128	256	512
64²					
time savings	74.4%	75.2%	73.0%	74.3%	73.7%
voxels retained	96.1%	96.9%	97.5%	97.5%	97.7%
voxel ratio	5.1%	4.3%	4.0%	4.0%	4.3%
128²					
time savings	78.4%	78.6%	78.0%	78.4%	79.5%
voxels retained	98.0%	98.4%	98.7%	98.7%	98.9%
voxel ratio	2.6%	2.1%	1.8%	1.8%	1.8%
256²					
time savings	80.6%	80.4%	79.8%	80.5%	81.2%
voxels retained	98.9%	99.2%	99.3%	99.3%	99.4%
voxel ratio	1.3%	1.1%	0.9%	0.9%	0.8%

Table 5: Vessel Data Results - Percentages

Size	Nodes				
	32	64	128	256	512
64²					
time savings	77.4%	76.7%	77.1%	79.1%	81.0%
voxels retained	98.2%	98.2%	98.7%	98.9%	98.9%
voxel ratio	2.8%	2.8%	2.2%	2.5%	2.1%
128²					
time savings	79.8%	79.7%	78.7%	79.8%	80.2%
voxels retained	98.9%	98.9%	99.4%	99.6%	99.6%
voxel ratio	1.7%	1.7%	1.1%	1.0%	1.0%
256²					
time savings	79.9%	80.5%	80.2%	80.1%	79.9%
voxels retained	99.4%	99.4%	99.7%	99.8%	99.8%
voxel ratio	0.9%	0.9%	0.6%	0.4%	0.4%

The evaluated voxel ratio of the right to left-eye is quite small. In the case of the head data with many terminated rays, the right-eye had to calculate 5% of the voxels evaluated in the left-eye. In the other data sets there are fewer terminated rays, and correspondingly fewer voxels evaluated in the right-eye, to as little as 0.2%. Figure 6 visually illustrates the limited number of samples evaluated only for the right-eye.

Table 6: Vorticity Data Results - Percentages

Size	Nodes				
	32	64	128	256	512
64²					
time savings	77.1%	78.5%	77.8%	80.7%	81.0%
voxels retained	99.1%	99.4%	99.3%	99.3%	99.1%
voxel ratio	1.2%	1.1%	1.3%	1.3%	1.9%
128²					
time savings	80.0%	80.4%	79.8%	80.1%	80.4%
voxels retained	99.4%	99.7%	99.8%	99.8%	99.8%
voxel ratio	0.8%	0.5%	0.5%	0.5%	0.6%
256²					
time savings	80.7%	80.8%	80.6%	80.7%	80.9%
voxels retained	99.6%	99.8%	99.9%	99.9%	99.9%
voxel ratio	0.6%	0.3%	0.2%	0.2%	0.2%

The evaluated voxel ratio increases inversely with image size, as explained above. A larger number of nodes will also cause the ratio to fall. As fewer rays terminate, fewer voxels will be evaluated only for the right-eye; there are more samples being reused from the left-eye view. Likewise, dense blocks which cause the retained voxel percentage to fall will increase the evaluated voxel percentage, as seen in Table 6.

Figure 6: Right-eye view of the vorticity data, with samples evaluated only for the right-eye illuminated.

Notice that the total percentage of voxels seen in the right-eye view (voxel retained percentage + voxel ratio percentage) is greater than 100% of the left-eye voxels. This is

primarily due to the volume bounding box, which is calculated for the left-eye. As seen in Figure 2, the left-eye rays traverse this box orthogonally, while the right-eye rays pass through at an angle, which is a longer traversal. Additionally, if the volume being rendered is longer in the z dimension (after the viewing transformation) than it is in the x dimension, there will be a greater number of right-eye rays passing through the volume than left-eye rays. The total number of voxels used in the right-eye view is usually about 1% greater than those in the left-eye view, and this percentage will drop as the viewing angle decreases and with cubically-shaped volumes.

3.1 Stereo Pair Generation and Other Volume Rendering Techniques

It is possible to use stereo pair coherence with other volume rendering techniques. For example, the dividing cubes method [CLLC88] produces point samples of non-transparent data which are then projected to the screen. These samples can be projected to the right eye using stereoscopic techniques developed for non-volumetric ray-tracing [AdHo93], and likely achieve similar time savings of 80-90% . Some methods, like marching cubes [LoCl87], produce a polygonal isosurface which is then rendered using traditional polygon scan-conversion methods. Polygons can be clipped, back-face culled, and projected simultaneously using coherence techniques [ABCH91]. If Gouraud or Phong shading is desired, the interior of the polygons must be shaded separately, since projected polygons will have the differing widths (although identical heights). Tests with relatively large polygons (an average polygon covering 50 pixels) produced a time savings of 36%. This savings is significantly smaller than the ray-traced volume rendering, and most of the loss occurs in the interior polygon color calculation. Since polygons extracted from a surface tend to be very small, the interior of polygons will be small, and a much larger savings is expected.

Other methods for projecting voxels directly, such as splatting [West90], can use the coherence to quickly find the pixels in the right-eye view affected by the voxel. Savings should be similar to that reported in this paper.

4. CONCLUSIONS AND FUTURE WORK

Stereoscopic images are a viable option for giving unambiguous depth information to volumetric rendering. Such images display the structure of the data without the need of continuous movement, which is usually employed in scientific visualization. A technique has been demonstrated which will generate stereoscopic pairs of volumetric data, creating the

second half of the pair in a fraction of the time of generating the first half. In this implementation, the second half of the pair was generated 74 - 81% faster than full rendering, with the same quality of image. The technique is more efficient with less opaque data, with data rendered at a higher resolution, and rendering with a larger number of processors.

The technique need not be used with a parallel render. A serial render could also use this stereoscopic algorithm, much in the same way that each node of the CM-5 individually uses the algorithm on a small portion of the volume. Also, volumetric ray-tracing modifications such as adaptive screen sampling, adaptive ray sampling, ray templates, and space-leaping (described in [Yage93]) could be implemented without affecting the reprojection technique. Using these methods may affect the time savings somewhat, since all of the left-eye voxels would be affected while only the right-eye voxels evaluated solely for that view would benefit from the faster method. This needs to be tested, but a significant retention of savings is expected.

Since the actual number of voxels rendered in the right-eye is much lower than the time savings suggests, it is likely that the time savings would be greater with optimized code. Also, when using multiple processors, a load balancing solution may increase savings, and should at the very least eliminate the minor drops in savings that were seen in Table 6, caused by poor load balancing.

5. ACKNOWLEDGMENTS

The authors would like to acknowledge the use of computing resources located at the Advanced Computing Laboratory (ACL) of Los Alamos National Laboratory, Los Alamos, NM 87545. This work was performed under the auspices of the United States Department of Energy.

6. REFERENCES

- [AABH93] Adelson SJ, Allen J, Badre AN, Hodges LF and Lawrence A, "Performance Comparison of Multiple Image Depth and Shape Cues." International Journal of Human Computer Interaction, 5, 4 (November 1993), 347-360.

- [ABCH91] Adelson SJ, Bentley JB, Chong IS, Hodges LF, and Winograd J, "Simultaneous Generation of Stereoscopic Views." Computer Graphics Forum, 10, 1 (March 1991), 3-10.
- [Adel94a] Adelson SJ, "Stereoscopic Projections: Parallel Viewing Vectors, Rotations, and Shears." Los Alamos National Laboratory Technical Report LA-UR-94-0115 (January 1994).
- [Adel94b] Adelson SJ, "Using stereoscopic imaging in visualization applications." SPIE Proceedings 2178: Visual Data Exploration and Analysis, (February 1994). To appear.
- [AdHo93] Adelson SJ and Hodges LF, "Stereoscopic Ray-Tracing." The Visual Computer, 10, 3 (December 1993), 127-144.
- [CLLC88] Cline HE, Lorensen WE, Ludke S, Crawford CR and Teeter BC, "Two Algorithms for the Three-Dimensional Reconstruction of Tomograms." Medical Physics, 15, 3 (May / June 1988), 320-327.
- [HBCP94] Hsu J, Babbs CF, Chelberg DM, Pizlo Z and Delp EJ, "A study of the effectiveness of stereo imaging truth or dare: Is stereo viewing really better?" SPIE Proceedings 2177a: Stereoscopic Displays and Applications V, (February 1994). To appear.
- [Hodg92] Hodges LF, "Time-Multiplexed Stereoscopic Computer Graphics," IEEE Computer Graphics and Applications, 12, 2 (March 1992), 20-30.
- [Hodg93] Hodges LF. Unpublished research results relayed in personal conversation.
- [HoMc91] Hodges LF and McWhorter SM, "Stereoscopic Display for Design Visualization." Image Communication, 4, 1 (November 1991), 3-13.
- [HoMc93] Hodges LF and McAllister DF, "Computing Stereoscopic Views," Stereo Computer Graphics and Other True 3D Technologies, ed. DF McAllister. Princeton: Princeton University Press, 1993.
- [Levo88] Levoy M, "Display of Surfaces from Volume Data." IEEE Computer Graphics and Applications, 8, 3 (May 1988), 29-37.
- [Lips79] Lipscomb JS, "Three-Dimensional Cues for a Molecular Computer Graphics System." Ph.D. Thesis, Department of Computer Science, University of North Carolina, Chapel Hill, NC (1979).
- [LoCl87] Lorensen WE and Cline HE, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm." Computer Graphics, 21, 4 (July 1987), 163-169.
- [MPHK93] Ma K, Painter JS, Hansen CD and Krogh MF, "A Data Distributed, Parallel Algorithm for Ray-Traced Volume Rendering." ACM SIGGRAPH Proceedings of the 1993 Symposium on Parallel Rendering, (October 1993), 15-22.
- [PoDu84] Porter T and Duff T, "Compositing Digital Images." Computer Graphics, 18, 3 (July 1984), 253-259.

- [ReBS90] Reinhart WF, Beaton RJ and Snyder HL, "Comparison of Depth Cues for Relative Depth Judgments." SPIE Proceedings 1256: Stereoscopic Displays and Applications, (February 1990), 12-21.
- [Rein91] Reinhart WF, "Depth Cueing for Visual Search and Cursor Positioning." SPIE Proceedings 1457: Stereoscopic Displays and Applications II, (February 1991), 221-232.
- [SoMi91] Sollenberger RL and Milgram P, "A Comparative Study of Rotational and Stereoscopic Computer Graphics Depth Cues." Proceedings of the Human Factors Society 35th Annual Meeting, (October 1991), 1452-1456.
- [ThMC91] The Connection Machine CM-5 Technical Summary.
- [WaAB93] Ware C, Arthur K and Booth KS, "Fish Tank Virtual Reality." INTERCHI '93, (April 1993), 37-42.
- [West90] Westover L, "Footprint Evaluation for Volume Rendering." Computer Graphics, 24, 4 (August 1990), 367-376.
- [Yage93] Yagel R, "Volume Viewing: State of the Art Survey." ACM SIGGRAPH Course Notes #41: Volume Visualization, (August 1993), 102-129.
- [YeSi90] Yeh Y and Silverstein LD, "Visual Performance with Monoscopic and Stereoscopic Presentation of Identical Three-Dimensional Visual Tasks." SID Digest XXI, (May 1990), 359-362.

DATE

FILMED

9 / 8 / 94

END