Efficient Interactive Pre-integrated Volume Rendering

Heewon Kye[1,3], Helen Hong[2], and Yeong Gil Shin[1,3]

[1] School of Computer Science and Engineering, Seoul National University
{kuei, yshin}@cglab.snu.ac.kr
[2] School of Computer Science and Engineering, BK21: Information Technology, Seoul National University, San 56-1 Shinlim-dong Kwanak-gu, Seoul 151-742, Korea
hlhong@cse.snu.ac.kr
[3] INFINITT Co., Ltd., Taesuk Bld., 275-5 Yangjae-dong Seocho-gu, Seoul 137-934, Korea

Abstract. Pre-integrated volume rendering has become one of the most efficient and important techniques in three dimensional medical visualization. It can produce high-quality images with less sampling. However, two important issues have received little attention throughout the ongoing discussion of pre-integration: Skipping over empty-space and the size of lookup table for a transfer function. In this paper, we present a novel approach for empty-space skipping using the *overlapped-min-max block*. Additionally, we propose a new approximation technique to reduce the dependent texture size so that it decreases the size of texture memory and the update time. We demonstrate performance gain and decreasing memory consumption for typical renditions of volumetric data sets.

1 Introduction

Pre-integrated volume rendering is a technique for reconstructing the continuous volume rendering integral. Utilizing a pre-processed look-up table (called pre-integration table), this method not only eliminates a lot of artifacts but also reduces the sampling rate for rendering. However, since this method uses two consecutive sample values as an index for the pre-integration table which is constructed before rendering for a given classified function, conventional acceleration techniques such as empty space skipping or interactive classification methods are not applied as it is.

Skipping empty space has been extensively exploited to accelerate volume rendering. However, pre-integrated volume rendering samples two consecutive points as a line segment, previous empty-space skipping methods could not be directly applied. The pre-integration table is indexed by three integration parameters: two consecutive sample values and the distance between those samples. To accelerate the pre-integration step, Engel *et al.* reduced the dimensionality of the table from three to two by summing a constant sampling distance [1]. Even though they used a two-dimensional pre-integration table, it is still bulky when rendering high-precision data such as 12 bits-per-voxel data which is common in medical applications. A 12-bit image requires 256 times more memory and updating time than an 8-bit image.

In this paper, we present a novel data structure, called the *overlapped-min-max block* for applying empty-space scheme to the pre-integrated volume rendering, and a new approximation technique for reducing the dimensionality of the table from two to

V.S. Sunderam et al. (Eds.): ICCS 2005, LNCS 3516, pp. 834 – 837, 2005.

one. We implement them on recent consumer graphics hardware and on software-only shear-warp rendering [2] and ray-casting [3]. With our accelerations, the rendering and classification speed is much faster for medical datasets while maintaining the image quality.

2 Overlapped Min-max Block for Empty-Space Skipping

Traditional rendering methods sample a value at a point in three-dimensional space to get a color and opacity. If a block is entirely transparent, additional samplings in the block can be skipped. Pre-integrated volume rendering samples two points to get their color and opacity. Since two sampling points form a line segment, or a slab, all the blocks that intersect the line segment should be transparent for skipping the sampling process. As shown in Fig. 1, a line segment may intersect at most three blocks in two-dimensional representation. Retrieving information three-times from the lookup table degrades the rendering performance. In addition, there is an overhead to determine which blocks are transparent (there are two cases such as Fig. 1a and Fig. 1b).

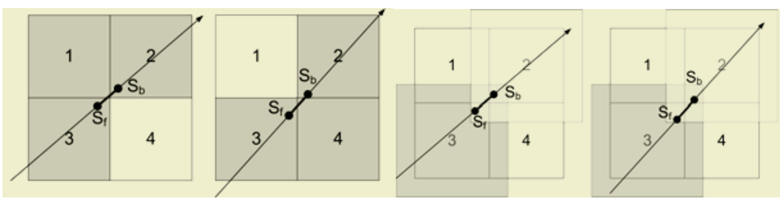

(a) (b) (c) (d)

Fig. 1. The overlapped min-max table for pre-integrated volume rendering. There are two sampling point S_f and S_b, and four blocks from *block1* to *block4*. S_f is in *block3* and S_b is in *block2*. To skip the line segment S_fS_b, (a) *block1*, *block2*, and *block3* have to be transparent, (b) *block2*, *block3*, and *block4* have to be transparent in the previous block structure. Moreover, there is overhead to determine the current situation is either (a) or (b). In our method, *the overlapped-min-max block*, each block holds some region jointly with its neighbors such as (c) and (d). To skip the line segment, we can test only one block, *block2* (or *block3*)

Making each line segment belonging to only single block, we can efficiently decide whether we skip or not by testing one block. For this, we modify the region covered by each block. Each block covers some region of which thickness is at least the sampling distance as shown in Fig. 1c and Fig. 1d. By overlapping the region of each block, we can easily test whether the block of a line segment is transparent using only that block. This scheme is especially efficient on graphics hardware for its simplicity.

3 Efficient Pre-integration Table

In accelerated pre-integrated rendering, opacity of i-th sample (α_i) is written as:

$$1-\alpha_i = 1-\alpha(s_f, s_b) \approx \exp\left(-\frac{1}{s_b - s_f}\left(T(s_b)-T(s_f)\right)\right) \tag{1}$$
$$where,\quad T(s) := \int_0^s \tau(s)ds.$$

Because equation (1) comprehends the ray segment integral of a transfer function, we do not need to consider the maximum of the Nyquiest frequencies of the transfer functions $\tau(s)$ with the scalar field s. Therefore, it is sufficient to sample a volume with relatively low frequency. The 2D lookup table to obtain α_i, requires a texture of which size is N^2, where N is density range. In order to generate color images, each texture entry requires 4 bytes. When an image has 12-bit pixel depth, the required texture size becomes 4096^2 and the required memory is 4096^2x4byte=64MB. The bigger lookup table brings the longer generation time and lower cache-efficiency. In this paper, we propose an 1D-texture lookup method that needs only 4096 entries. Because scalar values s are usually quantized equation (1) can be rewritten as:

$$1-\alpha_i \approx \sqrt[s_b-s_f]{\exp\left(-\sum_{s=s_f}^{s_b}\tau(s)\right)} = \sqrt[s_b-s_f]{\prod_{s=s_f}^{s_b}\exp(-\tau(s))} \approx \frac{1}{s_b-s_f}\sum_{s=s_f}^{s_b}\exp(-\tau(s)) = \left(S(s_b)-S(s_{f-1})\right)/\left(s_b-s_f\right)$$
$$where,\quad S(x)=\sum_{s=0}^{x}\exp(-\tau(s)),\ S(-1)=0. \qquad (2)$$

We modify a geometric average of transparency, into an arithmetic average as shown in equation (2). Therefore, only 1D-texture $S(x)$ needs to be stored to get the α_i. The color can be formulized in a similar fashion.

4 Experimental Results

Experiments have been performed on a PC equipped with Pentirum4 2.8GHz processor, 1GB main memory and an ATI 9800. Table 1 summarizes the performance using empty-space skipping (ESP). Obviously, skipping empty-space gains more performance. In hardware rendering, rendering time is reduced two times or more by using ESP. In software rendering, the gain of ESP is much bigger (10-30 times faster than w/o ESP). Since the angio dataset contains more empty space than the head dataset, the performance improvement factor of the angio dataset is bigger than that of the head dataset.

Table 1. A performance comparison in case of using empty-space skipping (msec)

Dataset	3D Hardware		Ray-Casting		Shear-Warp	
	with ESP	w/o ESP	with ESP	w/o ESP	with ESP	w/o ESP
BigHead	83.8	302.1	897.7	12409	747.6	7904
Angio	95.1	392.1	737.6	16183	339.6	9984

If a transfer function is fixed using 2D-texture (88.5ms) slightly faster than our 1D-texture (95.1ms) since 1D-texture requires two texture loads for classification while 2D-texture requires one. However, when the transfer function is changed, using our 1D-texture (101.2ms) is much faster than 2D-texture (824.4ms) because of lookup table creation time. Fig. 2 show rendered images of BigHead and Angio volume data, respectively. There is no noticeable difference between using 1D-texture and 2D-texture.

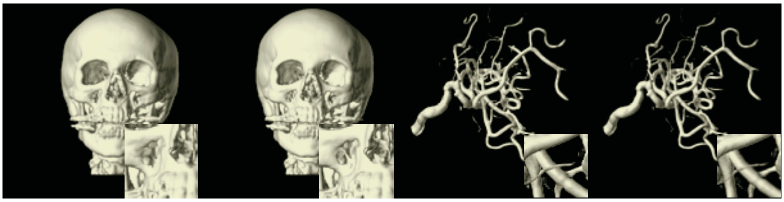

Fig. 2. The comparison of image quality for BigHead and Angio volume using (left) 1D-texture, (right) and 2D-texture

5 Conclusions

In this paper, we have proposed a new method to accelerate traversal and classification of both hardware and software based pre-integrated volume rendering. Using the overlapped-min-max block, empty-space skipping can be accomplished more efficiently and can be easily implemented in a hardware-based method. To reduce the classification time and memory consumption, a new approximation method of a lookup table is also proposed. With regard to image quality, we have presented the minimum bound of error theoretically. Experimental results show that our method produces the same quality of rendered images as the previous classification method of pre-integration.

References

1. K. Engel, M. Kraus, and T. Ertl. High-Quality Pre-Integrated Volume Rendering Using Hardware-Accelerated Pixel Shading. *Eurographics / SIGGRAPH Workshop on Graphics Hardware,* Los Angeles, CA, USA, August 2001.
2. P. Lacrout and M. Levoy, Fast Volume Rendering Using a Shear-Warp Factorization of the Viewing Transformation. *Proceedings of SIGGRAPH 94*, 451-458, 1994.
3. M. Levoy, Display of Surfaces from Volume Data, *IEEE Computer Graphics & Application*, 8: 29-37, 1988.