Available online at www.sciencedirect.com

SCIENCE DIRECT®

Future Generation Computer Systems 21 (2005) 573–581

FGCS
FUTURE GENERATION COMPUTER SYSTEMS

www.elsevier.com/locate/fgcs

Real-time rendering of 3D medical data sets

Kai Xie [a,*], Jie Yang [a], Y.M. Zhu [b]

[a] *Institute of Image Processing and Pattern Recognition, Shanghai Jiaotong University, 200030 Shanghai, China*
[b] *CREATIS, CNRS Research Unit 5515 and INSERM Unit 630, 69621 Villeurbanne, France*

Received 25 October 2004; received in revised form 12 December 2004; accepted 16 December 2004

Abstract

Interactive exploration of three dimension (3D) medical data sets is required by many applications, but the huge amount of computational time and storage space needed for rendering do not allow the visualization of large medical data sets by now. In this paper we present a new algorithm for rendering large medical data sets at interactive frame rates on standard PC hardware. The input data is converted into a compressed hierarchical wavelet representation in a preprocessing step. During rendering, the wavelet representation is decompressed on-the-fly and rendered using hardware texture mapping. The level of detail used for rendering is adapted to the local frequency spectrum of the data and its position relative to the viewer. Using a prototype implementation of the algorithm we were able to perform an interactive walkthrough of large medical data sets on a single of-the-shelf PC.

Keywords: Compression algorithms; Level of detail; Volume rendering; Wavelet

1. Introduction

In medical areas, we have to deal with large volumetric data sets that demand for an adequate visualization. An important visualization technique for the exploration of volumetric data sets is direct volume rendering: Each point in space is assigned a density for the emission and absorption of light and the volume renderer computes the light reaching the eye along viewing rays. The rendering can be implemented efficiently using texture mapping hardware: the volume is discretized into textured slices that are blended over each other using alpha blending [1].

Due to the enormous advances in graphics hardware, it is nowadays possible to perform this rendering technique in real-time on cheap of-the-shelf PCs [2]. However, the size of the data sets that can be processed is still very limited. Real-time rendering of large data sets is currently infeasible unless massive parallel hardware is used [3].

Most conventional hardware-texturing based volume rendering approaches are brute-force methods, requiring rendering time linear in the size of the data set.

* Corresponding author.
E-mail address: xie_kai2001@sjtu.edu.cn (K. Xie).

0167-739X/$ – see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2004.12.002

The rendering costs can be reduced dramatically by using a multi-resolution hierarchy. In this case, the rendering algorithm performs a projective classification to adapt the rendering resolution to the distance to the viewer, as proposed by LaMar et al. [4]. We will show formally in this paper that the rendering time for this technique is indeed O(log n) for a data set consisting of n^3 voxels. However, two problems still remain that prevent us from handling very large data sets: The first problem is the enormous size. Thus, the data must be stored out-of-core and swapped into main memory on demand. This leads to considerable bandwidth and latency problems. The second problem is the size of the voxel data that remains after projective classification: Although the size is O(log n), the constants in the "O-notation" are still much too high. The number of voxels exceeds by far the texture memory as well as the alpha-blending capacities of a commodity graphics board.

Our novel algorithm uses a hierarchical wavelet representation to tackle these problems: The volume is stored as a hierarchy of wavelet coefficients. Only the levels of detail necessary for display are decompressed and sent to the texturing hardware. The use of a wavelet representation allows us to compress the data by a ratio of typically 4:1 (lossless). During rendering, the wavelet representation allows us to analyze the local frequency spectrum in the data set and to adapt the rendering resolution to it. This way, we can reduce the size of the voxel set to be rendered considerably with no loss of image quality.

Using these techniques, we are able to render walkthroughs of large data sets in real time on a conventional PC. We will demonstrate an interactive walkthrough of the large data set at a resolution of 256^2 pixel, 20 frames per second and good image quality.

The remainder of the paper is structured as follows: In the next section, we will briefly review related work. Then, we will describe the hierarchical wavelet representation in Section 3. In Section 4, we will describe the rendering algorithm. Results are discussed in Section 5 and the conclusion in Section 6 with some ideas for future work.

2. Related work

Visualization of large volume data sets is a classical problem in medicine image visualization. In this section, we will give a brief overview of related work in the area of volume wavelet-based techniques, visualization algorithms and multi-resolution methods.

2.1. Wavelet based techniques

Wavelet based encoding has become a standard technique for 2D-image compression. The technique has been applied to the compression of volume data. The embedded zero-tree wavelet (EZW) [5] coding algorithm was introduced by Shapiro with excellent compression results. Later, Said and Pearlman proposed a more efficient coding algorithm using set partitioning in hierarchical tree (SPIHT), and implemented to both lossy [6] and lossless [7] compression of images. Kim and Pearlman [8] extended 2D to 3D for video, and Kim and Pearlman [9] utilized it for volume image compression. In this paper, we use Kim and Pearlman's lossless 3D SPIHT with asymmetric tree structure.

2.2. Volume visualization

The most efficient software-based technique for direct volume rendering is the shear-warp factorization by Lacroute and Levoy [10]. The technique can be adapted to exploit 2D-texturing hardware [11], achieving interactive frame rates. The usage of 3D texture mapping [12] allows for more flexibility and can provide a higher image quality. Recent visualization algorithms provide advanced shading techniques such as lighting [13], shadows [14], high quality post-classification using pre-integration technique [15], gradient magnitude modulation [16] or higher-dimensional transfer functions [17]. Our algorithm uses a pre-integration approach combining lighting and gradient magnitude modulation, as described in [18].

2.3. Multi-resolution rendering

The idea of multi-resolution volume rendering algorithms is to provide a spatial hierarchy to adapt the local rendering resolution to the projection onto the screen: An octree or a similar spatial data structure is build for the data set in a pre-processing step. Each node of the spatial hierarchy contains a certain part of