

Proceedings
**Computer Graphics
International 2003**


9-11 JULY 2003, TOKYO, JAPAN

TOKYO '92
CGI 2003 '83
Computer Graphics International

Proceedings

**Computer Graphics
International**

**Tokyo, Japan
July 9 – 11, 2003**



IEEE
**COMPUTER
SOCIETY**

<http://computer.org>

Los Alamitos, California

Washington • Brussels • Tokyo

Copyright © 2003 by The Institute of Electrical and
Electronics Engineers, Inc.
All rights reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number PR01946
ISBN 0-7695-1946-6
ISSN Number 1530-1052

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-732-981-0060
Fax: + 1-732-981-9667
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: reprints@computer.org

Editorial production by Bob Werner
Cover art production by Joe Daigle/Studio Productions
Printed in the United States of America by Applied Digital Imaging

IEEE
COMPUTER
SOCIETY



Preface

CGI (Computer Graphics International) was started by Professor Kunii in 1983 as CG Tokyo. After its initial success in Tokyo, the CGI conference were held in the international arena: Switzerland (1988, 1993, 2000), UK (1989, 1995, 2002), Singapore (1990), USA (1991), Australia (1994), Korea (1996), Belgium (1997), Germany (1998), Canada (1999), and Hong Kong (2001).

CGI 2003 returns to its origin, Tokyo, which is one of the most active cities in CG, VR, and multimedia in the world. It is the greatest pleasure for the local organizing committee to have the opportunity of organizing the memorial CGI2003 in Tokyo once again.

This year, we received a total of 74 submissions by December 9th, 2002, and we assigned at least three reviewers selected from the international program committee members for all of the papers. According to the scores given by the reviewers on January 30th, 2003, the local committee selected 38 regular papers and 18 short papers. Since the quality of the short papers were equally high to that of the regular papers, we decided to do a poster session with a short oral presentation for all the short papers. We are very pleased with the quality of these paper for CGI2003, and we must take this opportunity to thank all of the committee members and reviewers for doing an excellent job within the tight schedule.

The program committee is also grateful to the contributions of three invited speakers, Professor Nelson Max (Lawrence Livermore National Laboratory), Professor Hans-Peter Seidel (The Max-Planck-Institute Informatik), and Professor Hiroshi Iseki (Tokyo Womens Medical University). We would like to express our gratitude to all of the reviewers, program committee and local committee members, and sponsors who contributed to the great success of CGI 2003.

Finally, I would like to express my sincere thanks to Professor Nadia Magnenat-Thalmann (University of Geneva) and Professor T.S. Chua (The National University of Singapore) as sub-chair, and to Professor Tomoyuki Nishita (The University of Tokyo), Professor Yoshio Ohno (Keio University), Professor Katsumori Matsushima (The University of Tokyo) for their contribution to arrange this conference.

Masayuki Nakajima
Professor, Tokyo Institute of Technology

Table of Contents

Computer Graphics International — CGI 2003

Preface	ix
Session 1: CG Hardware	
Hardware Assisted Multichannel Volume Rendering	2
<i>A. Ghosh, P. Prabhu, A. Kaufman, and K. Mueller</i>	
Interactive Transfer Function Modification for Volume Rendering Using Compressed Sample Runs	8
<i>V. Srivastava, U. Chebrolu, and K. Mueller</i>	
Design and Implementation of Panoramic Movie System by Using Commodity 3D Graphics Hardware	14
<i>T. Yamamoto and M. Doi</i>	
Session 2: CG Modeling	
Interpolatory Subdivision Curves via Diffusion of Normals	22
<i>Y. Ohtake, A. Belyaev, and H. Seidel</i>	
Mesh Denoising via Iterative Alpha-Trimming and Nonlinear Diffusion of Normals with Automatic Thresholding	28
<i>H. Yagou, Y. Ohtake, and A. Belyaev</i>	
An Efficient Mesh Simplification Method with Feature Detection for Unstructured Meshes and Web Graphics	34
<i>B. Chen and T. Nishita</i>	
Invited Talk	
Clinical Application of Augmented Reality in Neurosurgical Field	44
<i>H. Iseki, Y. Muragaki, K. Naemura, M. Hayashi, T. Hori, and K. Takakura</i>	
Session 3: Natural Scene Modeling	
$1/f^\beta$ Noise-Based Real-Time Animation of Trees Swaying in Wind Fields	52
<i>S. Ota, M. Tamura, K. Fujita, T. Fujimoto, K. Muraoka, and N. Chiba</i>	
Modeling Lobed Leaves	60
<i>L. Mündermann, P. MacMurchy, J. Pivovarov, and P. Prusinkiewicz</i>	

Session 4: Rendering

- Data Compression for a Massive Image Data Set in IBMR _____ 68
C. Leung, P. Lam, and T. Wong
- Fast Volumetric Display of Natural Gaseous Phenomena _____ 74
S. Roettger and T. Ertl

Session 5: Virtual Reality

- Flexible Materials in Avango™ Virtual Environment Framework _____ 84
L. Nikitina, I. Nikitin, and S. Klimenko
- Immersive Vehicle Simulators for Prototyping, Training and Ergonomics _____ 90
*M. Kallmann, P. Lemoine, D. Thalmann, F. Cordier,
N. Magnenat-Thalmann, C. Ruspa, and S. Quattrocchio*
- VHD++Development Framework: Towards Extendible, Component Based VR/AR
Simulation Engine Featuring Advanced Virtual Character Technologies _____ 96
M. Ponder, G. Papagiannakis, T. Molet, N. Magnenat-Thalmann, and D. Thalmann
- Real Time Detection Interface for Walking on CAVE _____ 105
S. Barrera, P. Romanos, S. Saito, H. Takahashi, and M. Nakajima

Session 6: Texture

- Textile and Tile Pattern Design Automatic Cataloguing Using Detection of the
Plane Symmetry Group _____ 112
M. Valor, F. Albert, J. Gomis, and M. Contero
- Image Restoration Using Multiresolution Texture Synthesis and Image Inpainting _____ 120
H. Yamauchi, J. Haber, and H. Seidel
- Compressing Texture Coordinates with Selective Linear Predictions _____ 126
M. Isenburg and J. Snoeyink

Invited Talk

- Virtualizing Real-World Objects _____ 134
H. Lensch, J. Kautz, M. Goesele, J. Lang, and H. Seidel

Session 7: Non-Photo

- Rendering Artistic and Believable Trees for Cartoon Animation _____ 144
F. Di Fiore, W. Van Haevre, and F. Van Reeth
- “Nijimi” Rendering Algorithm for Creating Quality Black Ink Paintings _____ 152
Q. Guo and T. Kunii

Session 8: Application CG

Visualization of Distributed Processes Using “Data Jewelry Box” Algorithm _____	162
<i>Y. Yamaguchi and T. Itoh</i>	
Automatic Tracking of Face Sequences in MPEG Video _____	170
<i>Y. Zhao and T. Chua</i>	
A Framework for Transfer Colors Based on the Basic Color Categories _____	176
<i>Y. Chang, S. Saito, and M. Nakajima</i>	

Session 9: Real Time Rendering

Real-Time Visualization of Japanese Artcraft _____	184
<i>R. urikovi , R. Kimura, and K. Kolchin</i>	
A Real-Time Refraction Renderer for Volume Objects Using a Polygon-Rendering Scheme _____	190
<i>E. Ohbuchi</i>	
Fast Depth of Field Rendering with Surface Splatting _____	196
<i>J. K ivánek, J. Žára, and K. Bouatouch</i>	
Ray Tracing Height Fields _____	202
<i>H. Qu, F. Qiu, N. Zhang, A. Kauman, and M. Wan</i>	

Invited Talk

Volume Rendering for Curvilinear and Unstructured Grids _____	210
<i>N. Max, P. Williams, C. Silva, and R. Cook</i>	

Short Paper Session

Interactive Editing of Arbitrary Sub-Meshes _____	218
<i>F. Silva and A. Gomes</i>	
Cut-and-Paste Editing Based on Constrained B-spline Volume Fitting _____	222
<i>Y. Furukawa, H. Masuda, K. Miura, and H. Yamato</i>	
Sphairahedral Approach to Parameterize Visible Three Dimensional Quasi-Fuchsian Fractals _____	226
<i>K. Ahara and Y. Araki</i>	
A Method of Rendering Scenes Including Volumetric Objects Using Ray-Volume Buffers — Expanding to Render Scenes Including Overlapped Volumetric Objects _____	230
<i>K. Kajihara, H. Takahashi, and M. Nakajima</i>	
Time-Sequence Dynamic Virtual Images _____	236
<i>Z. Li, J. Sun, and J. Zhang</i>	
Using Virtual Reality to bring Singaporean Heritage to Life _____	240
<i>M. Song, T. Elias, W. Müller-Wittig, and T. Chan</i>	
Modeling Rocky Scenery Taking into Account Joints _____	244
<i>T. Ito, T. Fujimoto, K. Muraoka, and N. Chiba</i>	

Consistent Visualization and Querying of GIS Databases by a Location-Aware Mobile Agent _____	248
<i>S. Lodha, N. Faaland, G. Wong, A. Charaniya, S. Ramalingam, and A. Keller</i>	
Robust Watermarking of 3D Polygonal Models Based on Vertice Scrambling _____	254
<i>Y. Zhi-qiang, H. Ip, and L. Kwok</i>	
A Framework to Customize a Face Model for Reusing Animation _____	258
<i>K. Zhang, Z. Huang, and T. Chua</i>	
Camera Self-Calibration Using the Geometric Structure in Real Scenes _____	262
<i>P. Liu, J. Shi, J. Zhou, and L. Jiang</i>	
An Inverse Kinematics Method for 3D Figures with Motion Data _____	266
<i>T. Komura, A. Kuroda, S. Kudoh, T. Lan, and Y. Shinagawa</i>	
Natural Human Animation via Learning with Dynamic Manipulability _____	272
<i>T. Mukai, S. Kuriyama, and T. Kaneko</i>	
Cartoon Blur: Non-Photorealistic Motion Blur _____	276
<i>Y. Kawagishi, K. Hatsuyama, and K. Kondo</i>	
A Production Method of Reusing Existing 2D Animation Sequences _____	282
<i>F. Sumi and M. Nakajima</i>	
An Educational Tool for Basic Techniques in Beginner's Pencil Drawing _____	288
<i>S. Takagi, N. Matsuda, M. Soga, H. Taki, T. Shima, and F. Yoshimoto</i>	
Generation of Varing Line Thickness _____	294
<i>S. Saito, A. Kani, Y. Chang, and M. Nakajima</i>	
Visualization of Musical Pitch _____	300
<i>P. McLeod and G. Wyvill</i>	
Session 10: Animation	
Extraction of Motion Individuality in Sports and its Application to Motion of Characters with Different Figures _____	306
<i>R. Kawasaki, Y. Kitamura, and F. Kishino</i>	
Creating Various Styles of Animations Using Example-Based Filtering _____	312
<i>R. Hashimoto, H. Johan, and T. Nishita</i>	
Diorama Engine — A 3D Directing Tool for 3D Computer Animation Production _____	318
<i>K. Mikami and T. Tokuhara</i>	
Reaching Volumes Generated by Means of Octal Trees and Cartesian Constraints _____	324
<i>I. Rodríguez, M. Peinado, R. Boulic, and D. Mezjat</i>	
Color Plates _____	331
Author Index _____	337

Session 1

CG Hardware



Hardware Assisted Multichannel Volume Rendering

Abhijeet Ghosh* Poojan Prabhu[†] Arie E. Kaufman* Klaus Mueller*

Center for Visual Computing (CVC) and Computer Science Department
Stony Brook University
Stony Brook, NY 11794-4400, USA

Abstract

We explore real time volume rendering of multichannel data for volumes with color and multi-modal information. We demonstrate volume rendering of the Visible Human Male color dataset and photo-realistic rendering of voxelized terrains, and achieve high quality visualizations. We render multi-modal volumes utilizing hardware programmability for accumulation level mixing, and use CT and MRI information as examples. We also use multi-board parallel/distributed rendering schemes for large datasets and investigate scalability issues. We employ the VolumePro 1000 for real time multichannel volume rendering. Our approach, however, is not hardware-specific and can use commodity texture hardware instead.

Keywords: Multichannel, post-classification, transfer functions, multi-modal, volume mixing, Visible Human, terrain rendering, parallel rendering, distributed rendering, image compositing, VolumePro 1000, graphics hardware.

1. Introduction

Many volumetric datasets in medical and scientific applications have multichannel information. Examples include the Visible Human Male (VHM) dataset from the National Library of Medicine, which has CT, MRI as well as color information from photographs. Other examples are aerial or satellite color ortho-photographs of terrains (used in height-field visualization systems), and multi-spectral (remote-sensing) satellite data. This type of data offers exciting possibilities for photo-realistic volume visualization as well as better understanding of the data due to the multichannel information. In this paper, we explore hardware assisted volume rendering of multichannel data with examples taken from the VHM and terrain rendering applications.

Direct volume rendering is expensive in terms of computation and memory in comparison with surface based ren-

dering. Volume rendering is preferred for medical visualization because of its ability to explore the internal structures such as bones and tissues. Polygon based rendering has been traditionally the choice for terrain visualization as terrains are mostly flat surfaces which can be modelled well with polygons and real time rendering can be achieved with texture mapping hardware using various LOD techniques [7]. Volume rendering still has many advantages for terrain visualization as the elevation maps can be easily voxelized into very high-resolution 3D volumes. Also, the voxel-based model is scene complexity independent, and offers a better representation for amorphous phenomena such as clouds, haze and fire. Therefore, volume rendering has been successfully applied to terrain visualization [3, 18].

Interactive rendering speed is required for the ease of exploration of large volumetric datasets and had been restricted for some time to high-end graphics workstations, due to the requirement of trilinear interpolations for image quality. Direct volume rendering with commodity 2D texture mapping [1, 5, 16] as well as 3D texture mapping have been successfully exploited for volume rendering. However, most PC platform texture mapping hardware have a limited on-board texture memory (typically 64-128MB) and slow texture-memory access rates, especially for 3D textures, limiting the size of volumes that can be rendered interactively in a single pass.

The above factors make special purpose hardware attractive for direct volume rendering of large volumes. The special purpose volume rendering hardware VolumePro 500 [13], which evolved out of our Cube-4 architecture [14], supports real time frame rates (30 frames/sec) for a 256³ volume. The second generation hardware, the VolumePro 1000, is able to associate multiple channels of information (vector) with a voxel, not just a scalar density. It also has the capability of rendering super-volumes (i.e., volumes larger than the on-board memory) in real time with multi-pass rendering, and its pipeline is programmable, similar to commodity hardware. Given that this hardware is in its initial years of development, we hope that with future developments, the price/performance ratio will be more attractive.

* {abhijeet, poojan, ari, mueller}@cs.sunysb.edu

[†] now at Frounhofer Institute of Computer Graphics, Germany

In this paper, we demonstrate a post-classified single pass volume rendering of multichannel data (e.g., $RGB\alpha$ color volumes and multi-modal rendering of CT and MRI data) with graphics hardware, and multi-board real-time rendering for large datasets. We employ the VolumePro 1000 for real time rendering of multi channel data. However, our approach is not specific to VolumePro and can be implemented on commodity texture hardware as well.

Section 2 discusses the modeling and rendering of the VHM color volume and issues with transfer functions and gradients for multichannel data. We also explore multi-modal visualization of the CT and MRI datasets of the VHM in Section 3. Multi-board rendering for large datasets is discussed in Section 4. In Section 5, we briefly discuss the rendering pipeline of the VolumePro 1000 and some implementation issues. Alternative texture mapping hardware rendering schemes are also proposed. Finally, we discuss realistic terrain modeling and navigation as an application in Section 6.

2. Multichannel Color Volume Rendering

Color photographic volumes greatly simplify the task of creating realistic volume rendered images as the appropriate color for each voxel is already known from the photographs. Photographic volumes however offer a new challenge: determining the opacity for each voxel in the dataset. In traditional volume rendering, the design of effective color (1D to 3D mapping) and opacity (1D to 1D mapping) transfer functions for meaningful exploration of the data has been challenging [15]. In contrast, volume rendering from photographic data requires an opacity transfer function from the vector color fields (3D to 1D) and which is complicated by the non-linear nature of color spaces.

A way to solve this problem is to associate opacity values for the color voxels from an auxiliary density volume, such as CT, if available. There are two problems with this approach: a modality such as CT is good for external boundaries but does not capture internal details such as soft tissues well and the density volume has to be registered with the photographic volume. Registration of a large volume such as the VHM is difficult to automate.

Hence, we approach the direct volume rendering of the VHM color volume from the photographic information itself. We use the CIE $L^* u^* v^*$ color space to obtain a perceptually uniform representation of the color volume [4]. The L^* component corresponds to the linear lightness of color, and hence has high values for bright areas of the volume (e.g., bones, skin and light colored organs such as the brain) (see Fig. 1a). The u^* component captures the chromatic changes in the red-green colors. Hence, it is sensitive to changes in the “redness” of tissues (e.g., muscle to bone) (see Fig. 1b). Our hardware rendering of the VHM is able




Figure 1. Multichannel volume rendering: (a) L^* color component used as the density. (b) u^* color component used as the density.

to accurately simulate these properties of the color space. Currently, graphics hardware can store information in up to four fields, typically meant for $RGB\alpha$. We assign the RGB information from the slices and the density information from the corresponding L^* or u^* color component that we calculate in a pre-process.

We propose a way to incorporate color difference gradients for shading in hardware rendering. Gradients can be calculated in hardware based on the voxel information. Hardware programmability allows us to access information of the neighboring voxels, which we can use for various kinds of gradient calculation. The changes in color values along X , Y and Z , expressed as the triple ($grad.x$, $grad.y$, $grad.z$), is the color difference gradient vector. We implement color difference gradient magnitude calculation with hardware (see Fig. 2a) by using the difference of the information in the Red channel across a voxel along the X direction as $grad.x$, and likewise for the other two color channels. The perceptual difference between two colors is to a good approximation proportional to the Euclidean distance between them. This is specially true for the CIELUV space. For color difference gradients in the CIELUV space, we sacrifice the true color volume rendering due to the limited number of color channels and render only with pseudo colors from an $RGB\alpha$ transfer function (see Fig. 2b).

We also incorporate gradient boundary enhancement by using first and second order directional derivatives along the gradient direction [11]. The first derivative in the direction of the gradient is the gradient magnitude itself and the second directional derivative can be well approximated by the gradient of the gradient magnitude. The first directional derivative can be currently calculated in hardware rendering. For the second directional derivative, we pre-calculate the first directional derivative and store them in separate channels. Then, using the existing algorithm, calculate the




Figure 2. Multichannel volume rendering with hardware assisted gradients: (a) Color difference gradients of RGB color space. (b) Gradients of LUV space with second derivative along gradient direction.

gradient based on the information in these channels, resulting in second directional derivative on the fly (see Fig. 2b).

3. Multi-Modal Volume Rendering

In this section, we describe real-time multi-modal volume rendering. Various levels (image, accumulation and illumination) of volume mixing and rendering pipelines have been proposed for multi-modal rendering [2]. Image level intermixing is not efficient as it requires one rendering pass per volume and an additional merging, and also does not produce good quality visualization. Though accumulation level mixing produces good quality visualization, it is a slow process (mixing happens on a per sample basis along the ray) and our aim is to accelerate this using hardware. In our scheme, each modality is associated with a color channel for single-pass rendering in hardware. This scheme also ensures a color coding scheme for information from the different modalities. For the purpose of demonstration, we use information from two kinds of volumes, CT and MRI of the VHM dataset. However, our scheme is generic enough to be used for other multi-modal datasets.

A simple but lossy scheme for mixing data from two modalities such as CT and MRI is to sample from both using a density threshold, assigning value to the voxel from CT if the corresponding CT density is above the threshold (high density for bones and surfaces) and from MRI otherwise (low density for soft tissues). Hence, we get the best of both modalities. We use red color for MRI (density information replicated to Red and α channels) and gray for CT (density information replicated to all four $RGB\alpha$ channels) (see Fig. 3a) as our color coding scheme.




Figure 3. (a) Simple multi-modal volume mixing of CT and MRI data by sampling based on density threshold. Mixing with arithmetic operation: (b) gradients specified on CT while MRI is rendered as a point cloud, and (c) rendering with gradient modulation.

The above approach has the drawback of information being lost due to the thresholding during sampling. A better way to implement mixing would be to let the user select the threshold while exploring the data. We achieve this by employing the arithmetic and logical operation supported in hardware. We store the CT and MRI information in separate fields (color channels in hardware) of the voxel. Then, we compute CT - MRI and MRI - CT simultaneously and output the results of this Diff operation (with negative values clamped) on two separate color channels (green and blue in our case). The result is that for every voxel, only the greater of the two components of the voxel contributes to the final accumulation along the ray. This scheme is not lossy and the segmentation can be dynamically altered using transfer functions for data exploration. One issue that arises is that gradients need to be calculated across the different modalities separately since they represent different volumes and this may not be supported in hardware. Hence, we specify gradients to be calculated only from one modality (voxel field) at a time, and not on the entire (multi-modal) voxel (see Fig. 3b). The above mixing scheme is used for the color channels only (for color coding), while we use $\text{Max}(\text{CT}, \text{MRI})$ for the α channel for well defined iso-surfaces using gradient modulation (see Fig. 3c).

Another approach to mixing is inclusive opacity, where all the volumes contribute to the voxel final color and/or opacity. A convenient way to identify voxels overlapped by both CT and MRI is by rendering them using a third orthogonal color (red in our case). We render the multi-modal volume with the function $1 - (\text{CT} - \text{MRI}) - (\text{MRI} - \text{CT})$ for the red channel corresponding to voxels with information from both modalities. We still use the Diff operator for the green and blue channels and the Max operator for the α channel (see Fig. 4). More sophisticated operations such as conditionals should be soon available in graphics hardware adding flexibility to the volume mixing schemes.




Figure 4. Multi-modal volume rendering with inclusive opacity: (a) Voxels with information from both CT and MRI; (b) Voxels with more CT.

4. Multi-Board Rendering

We utilize multiple units for rendering large volumetric datasets interactively. Many such visualization systems exist, ranging from those using multiple CPUs and graphics pipes in parallel [6] to cluster-based systems using special-purpose hardware such as VolumePro 500 [8, 12]. Here, we investigate the efficient parallelization of multiple graphics hardware boards on the PC platform.

We implement image-partitioned rendering by loading the entire volume on all the available boards, but restricting the range of the image and depth buffers to be filled by each board. This way, we distribute the rendering task to the boards uniformly. A limitation of this approach is that the size of the volume that can be rendered is limited by the memory size of the each board. Therefore, super-volumes have to be rendered using object-partitioned parallelism. A super-volume is divided into multiple sub-volumes, each of which can fit on one on-board memory. Each board renders its sub-volume and the resulting images have to be composited in order of depth. Hence, object-ordered parallelism seems more useful for large datasets.

We investigated distributed rendering on multiple PCs with one rendering engine per node as PC architecture does not commonly support multiple graphics boards. We setup one node as the master (control) and the others as rendering slaves. The master distributes the viewing parameters per frame to the slaves and collects the rendered images from them for compositing and display on a per frame basis. The slaves render their corresponding sub-volume after receiving the viewing parameters and send the rendered image to the master once every frame. This setup requires a high bandwidth low latency interconnect and network interface technology, such as Myrinet or Gigabit Ethernet, to support the sustained heavy network traffic.

Image compositing of rendered sub-volumes can be a computationally expensive task as the compositing opera-




Figure 5. Multi-board volume rendering: (a)-(b) rendered sub-volumes; and (c) image composited on texture mapping hardware

tor has to be applied over every pixel of the image. We utilize the register combiner units on the texture mapping hardware (GeForce2) to carry out this per-pixel compositing and display at the master node in an efficient hardware accelerated fashion (see Fig. 5). The available sustained network bandwidth as well as the loading of the image data into texture memory are the overhead in this distributed rendering scheme and cost a few frames in performance.

We can also use multiple boards to cache large volumes ahead of rendering and at run-time only render the desired sub-volume. This is useful for terrain visualization (see Section 6) where the terrain may be large but only portions are visible at any point. Caching has been utilized for walk-throughs for reusing previously rendered images in subsequent frames [17]. We divide the terrain volume into slabs that can fit in the on-board memory. This way, we maintain real time frame rates for a fly-through by rendering only relatively small slabs at a time. This scheme also helps overcome the loading-unloading latency during fly-through as all required slabs are already on-board.

5. Implementation with VolumePro 1000

VolumePro 1000 supports real-time volume rendering of up to 513^3 volumes. The hardware performs the following basic operations on data at voxel or sample points: (1) Gradient estimation (central difference), (2) Classification/Interpolation (trilinear), (3) Illumination, and (4) Compositing. Alpha correction, accumulation, and early ray termination are included in the compositing process. The voxels in VolumePro memory are composed of up to a maximum of 4 fields. The classification function of VolumePro consists of transfer functions for color and α lookup and arithmetic and logical units.

We implement multichannel volume rendering with VolumePro 1000, utilizing calls of the Volume Library Interface (VLI) (C++ API for VolumePro). We render the color and multi-modal volumes on the VolumePro by assigning 8 bits per color channel or modality to a separate field of

the voxel. These fields can also be used to tag segmented data. Mapping of the voxel fields through transfer function LUTs is supported in VolumePro. This can be implemented with commodity texture hardware using multi-texturing and dependent texture lookups. The ALU units of VolumePro provide additional programmability, similar to the register combiner units of commodity graphics hardware.

We down-sampled the original VHM photographs from a resolution of 2048×1216 pixels to 512×256 pixels. Our color space conversion was performed by first converting RGB to CIE XYZ space using the *XYZitu601-1 (D65)* standard conversion matrix, and then converting to CIE $L^*u^*v^*$ color space.

For multi-modal volume rendering, we modulated each voxel field with a separate transfer function and used the ALU units for implementing the Diff and Max operation for volume mixing. CT slices from the VHM are of 512×512 resolution and the MRI slices are 256×256 . Hence, we first interpolate the MRI slices to 512×512 resolution. Also, the MRI data is available for every fourth slice of the CT. Therefore, we interpolate between slices with a large filter support to smooth out the staircase effect.

Though we experimented with multi-board VolumePro 1000 rendering on a single machine, we discovered it to be inefficient due to the PC architecture. We solve this issue by switching to distributed rendering with object ordered parallelism. Our test bed is a dual node PC cluster with one VolumePro per node. The nodes are connected with a high-speed 1 Gbps Myrinet point-to-point connection and we use the low level API GM for packet transfer over the Myrinet.

6. Applications and Results

We create realistic visualization of volumetric terrain as another application of multichannel color volume rendering on VolumePro 1000. One of the primary sources of data in terrain visualization is the terrain height field and the corresponding texture of the terrain. The latter is typically obtained from an aerial or satellite ortho-photograph. We construct the terrain dataset by first voxelizing the height-field data, creating an iso-surface of voxels. We scale all heights in our height-field data to a maximum height of 64 voxels as terrains do not generally have sharp variations in altitude. The terrain dataset is constructed as an $RGB\alpha$ volume, with one channel per field. At location (x, y, z) , we obtain RGB information from the corresponding (x, z) location of the terrain ortho-photograph (texture) (see Fig. 6a). For the α channel we assign a constant high value in order to make the surface opaque.

Finally, we utilize VolumePro 1000 ability to embed polygonal objects within volumes to create the visual effect of an F-15 flying over volumetric terrain (see Fig. 6b). Though great for visualization, this is a costly process. We




Figure 6. (a) Original ortho-photograph (512×512) of the terrain; (b) High quality multichannel volume rendering of terrain ($1024 \times 64 \times 1024$) with embedded polygonal F-15.




Figure 7. Multichannel volume rendering of volumetric terrain with amorphous clouds

also produce high quality rendering of terrain with amorphous phenomenon such as clouds (see Fig. 7) and achieve interactivity with hardware volume rendering.

Our experiments are conducted using VolumePro 1000 boards with 1GB of on-board memory. Some of this memory is allocated for on-board image and depth buffers. Hence, we restrict the largest 32-bit $RGB\alpha$ color volume in a single pass to 950MB. For the visible Human Male, this evaluates to 512×256 resolution per slice and a total of 1871 slices. We divide the terrain volume into slabs, and restrict the slab size to $1024 \times 64 \times 1024$. We cache three such slabs on one board at a time for a total volume size of $3072 \times 64 \times 1024$.

For our parallel rendering experiments, we have experimented with up to four boards in parallel PCI (33MHz-32bit) slots of a Pentium III 1GHz machine, as well as in parallel PCI (66MHz-64bit) slots of an Athlon 1.2 GHz machine (see Table 1). Higher bandwidth of the 66MHz-64bit PCI bus gives consistently better performance. Unfortunately, the PCI slots on these machines share the same PCI bus and hence the performance does not scale well with the number of boards. Our distributed rendering experi-

Table 1. Frame rates for multi-board parallel rendering of VHM on one PC

PCI bus	Single Board	2 Boards on 1 PC	4 Boards on 1 PC
33MHz-32bit	8	14	18
66MHz-64bit	14	19	24

ment with VHM was configured using two nodes each with $512 \times 256 \times 936$ sub-volume and one or two boards. The frame rates for a 66MHz-64bit PCI bus are 25 for 1 board and 33 for 2 boards per node. Tiling smaller slabs of size $1024 \times 64 \times 1024$ within each board for volumetric terrain visualization provides real-time performance of 32 and 48 frames/sec for 32MHz-32bit and 66MHz-64bit PCI bus, respectively.

7. Conclusions and Future Work

We have demonstrated high-quality visualization of multichannel data with volume rendering through proper modeling of the data and hardware assisted single pass rendering. We achieve various mixing schemes for single pass multi-modal volume rendering using hardware units. Interactive speed has been achieved for super-volumes using multi-board rendering. Though we have used special purpose VolumePro 1000, our approach is generic and can be realized on commodity texture hardware too.

We are in the process of setting up a visualization PC cluster, initially with 12 rendering nodes and a display client. One issue that we need to resolve in this cluster is that of fast image composition, maybe using the efficient binary swap composition strategy [9]. We intend to overcome the slow reads of image data from the rendering unit using special-purpose compositing hardware [10, 12], in order to match the compositing speed with that of the rendering.

Acknowledgements

This work is supported by ONR grant N000140110034. We thank Andy Vesper and TeraRecon for information on the VolumePro 1000, and thank Kevin Kreeger, Wei Li and Huamin Qu for valuable discussions.

References

[1] B. Cabral, N. Cam, and J. Foran. Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *Proc. ACM Symp. on Volume Visualization*, pages 91–98, 1994.

[2] W. Cai and G. Sakas. Data intermixing and multi-volume rendering. *Computer Graphics Forum*, 18(3):359–368, September 1999.

[3] D. Cohen-Or, E. Rich, U. Lerner, and V. Shenkar. A real-time photo-realistic visual flythrough. *IEEE Trans. on Visu-*

alization and Computer Graphics, 2(3):255–265, September 1996.

[4] D. Ebert, C. Morris, P. Rheingans, and T. Yoo. Designing effective transfer functions for volume rendering from photographic volumes. *IEEE Trans. on Visualization and Computer Graphics*, 8(2):183–197, April-June 2002.

[5] K. Engel, M. Kraus, and T. Ertl. High quality pre-intergrated volume rendering using hardware-accelerated pixel shading. In *Proc. SIGGRAPH/Eurographics Workshop on Graphics Hardware 2001*, pages 9–16, 2001.

[6] J. Kniss, P. McCormick, A. McPherson, J. Ahrens, J. Painter, A. Keahey, and C. Hansen. Interactive texture-based volume rendering for large data sets. *IEEE Computer Graphics and Applications*, 24(4):52–61, July-August 2001.

[7] P. Lindstrom, D. Koller, W. Ribarsky, L. Hodges, N. Faust, and G. Turner. Real-time, continuous level of detail rendering of height fields. In *Proc. SIGGRAPH 96*, pages 109–118, August 1996.

[8] S. Lombeyda, L. Moll, M. Shand, D. Breen, and A. Heirich. Scalable interactive volume rendering using off-the-shelf components. In *Proc. Symp. on Parallel and Large-Data Visualization and Graphics*, pages 115–121, October 2001.

[9] K. Ma, J. Painter, C. Hansen, and M. Krogh. Parallel volume rendering using binary-swap compositing. *IEEE Computer Graphics and Applications*, 14(4):59–68, July-August 1994.

[10] L. Moll, A. Heirich, and M. Shand. Sepia: Scalable 3D compositing using PCI pamette. In *Proc. IEEE Symp. on Field Programmable Custom Computing Machines*, pages 146–155, April 1999.

[11] C. Morris and E. Ebert. Direct volume rendering of photographic volumes using multi-dimensional color-based transfer functions. In *EUROGRAPHICS - IEEE TCVG Symp. on Visualization*, pages 115–124, 2002.

[12] S. Muraki, M. Ogata, K. Ma, K. Koshizuka, K. Kajihara, X. Liu, Y. Nagano, and K. Shimokawa. Next generation supercomputing using PC clusters with volume graphics hardware devices. In *Proc. IEEE Supercomputing Conference*, November 2001.

[13] H. Pfister, J. Hardenbergh, J. Knittel, H. Lauer, and L. Seiler. The volumepro real-time ray-casting system. In *Proc. SIGGRAPH 99*, pages 251–260, August 1999.

[14] H. Pfister and A. Kaufman. Cube-4: A scalable architecture for real-time volume rendering. In *Proc. IEEE Symp. on Volume Visualization*, pages 47–54, October 1996.

[15] H. Pfister, B. Lorensen, C. Bajaj, G. Kindlmann, W. Shroeder, L. Avila, K. Martin, R. Machiraju, and J. Lee. The transfer function bakeoff. *IEEE Computer Graphics and Applications*, 21(3):16–22, May-June 2001.

[16] C. Rezk-Salama, K. Engel, M. Bauer, G. Griener, and T. Ertl. Interactive volume rendering on standard PC graphics hardware using multi-textures and multi-stage rasterization. In *Proc. SIGGRAPH/Eurographics Workshop on Graphics Hardware 2000*, pages 109–118, 2000.

[17] J. Shade, D. Lischinski, D. Salesin, T. DeRose, and J. Snyder. Hierarchical image caching for accelerated walkthroughs of complex environments. In *Proc. SIGGRAPH 96*, pages 75–82, August 1996.

[18] M. Wan, H. Qu, and A. Kaufman. Virtual flythrough over a voxel-based terrain. In *IEEE Virtual Reality Conference*, pages 53–60, 1999.

Interactive Transfer Function Modification For Volume Rendering Using Compressed Sample Runs

Vivek Srivastava Uday Chebrolu Klaus Mueller

Center for Visual Computing, Computer Science, Stony Brook University

Abstract

This paper describes a software-based method for interactive transfer function modification. Our approach exploits the fact that, in general, a user will rarely want to modify the viewpoint and the transfer functions at the same time. In that spirit, we optimize the latter by first fixing the viewpoint and then storing a compressed list of samples along each ray. Then, each time the transfer function is modified, the algorithm traverses the compressed sample lists, decompressing the runs and compositing the newly colored samples along each ray until full opacity is reached. Since the expensive sample interpolation and shading is no longer necessary, we can obtain near-interactive framerates for a variety of datasets, while our RLE-based compression of linearly varying sample runs helps keep the storage complexity down, with little decompression overhead. Decompression cost is reduced by storing decompression results in an on-the-fly constructed 2D table.

1. Introduction

A volumetric data object is described as a space-filling three-dimensional grid of discrete sample points, which, in turn, support the interpolation of any arbitrary point within the grid's 3D bounding box. A great variety of disciplines generate, use, and modify volumetric data. Examples are the medical field in diagnosis and surgical simulation, engineering in CAD/CAM prototyping, the oil and gas industry in natural resource exploration, designers and artists in virtual sculpting and industrial design, educators in teaching biology, chemistry, and anatomy, the computer game industry in the generation of realistic natural phenomena, computational scientists in scientific data exploration, and the business world in visual data mining.

Volume rendering is the process of exploring the volumetric data using visuals. The exploration process aims to discover and emphasize interesting structures and phenomena embedded in the data, while de-emphasizing or completely culling away occluding structures currently not of interest. In volume rendering, two main instruments exist that control the exploration process: view navigation and transfer functions. Both are essential. The former determines the spatial position and orientation from which the user observes the scene. The latter controls the look-and-feel of the scene itself, which is done by ways of the transfer

functions that map the raw object density data to color and transparencies. This also constitutes a navigation task, performed in a 4D transfer function space, assuming 3 axes for RGB color and one for transparency (or opacity). Note that while spatial navigation also exists in the surface rendering of polygonal objects, transfer function navigation is unique to volume rendering.

Transfer function space navigation can be a time consuming task for two main reasons: (i) the space to be explored is large, and (ii) the volume rendering is, at least at the present time, only interactive when supported by hardware [2][5][17][21][23]. Hardware solutions, however, are only practical for moderate dataset sizes, unless expensive machines are used [20]. On the other hand, interactive software solutions also exist, and a prominent representative for these is the shear-warp algorithm [12], which, however, has a considerable quality-speed trade-off, especially at odd viewing angles [25].

Addressing the complexities associated with the magnitude of the transfer function space, a number of techniques have been devised that use gradient histograms and other derived data to point out interesting locations in transfer function space [1][10][11][26]. The user may then use these hints to perform a more detailed search around these critical points. Another suggested strategy is to render a large number of images with arbitrary transfer function settings and present these to the user, who then selects a subset of these for further refinement by ways of genetic algorithms. This approach is taken by the Design Galleries project [16], which is based, in part, on the method published in [8]. A good sample of the existing approaches were squared off in a recent panel, termed the 'Transfer Function Bake-off', at the Visualization 2000 conference [22].

Common to all transfer function exploration efforts is the need for rendering the volume using the new settings. A key observation here is that the user rarely changes the viewing parameters *and* the transfer function simultaneously. Usually, the object stays fixed in space, and only the transfer function is modified, or vice versa, with the reason for this being that the user is simply too occupied with one of the two tasks to deal with the other. While view navigation can use temporal coherencies to speed up the rendering [3][19][27], for this research we sought to identify the coherencies that exist in the task of transfer function navigation. We start by realizing that the main work in volume rendering is formed by the following pipeline: interpolation of the ray samples, illumination (shading), coloring, and com-

Email: {vivek, uday, mueller}@cs.sunysb.edu

positing. Since we leave both the view and the light source position unchanged, there is no need to perform the sample interpolation and shading for every new transfer function setting. Instead, we may cache away the densities and shades of the samples for repeated use in the transfer function exploration process. This was recognized quite some time ago in [15]. A caveat with this approach, however, is that all samples along a ray, within the limits of the volume's bounding box, are now potentially visible, depending on the user's choice of transfer function setting. Hence, they must all be cached, which gives rise to a storage complexity approaching that of the volume. A suitable compression scheme is clearly needed, with the requirement of fast decompression speed since our main goal is to achieve an interactive system for transfer function modification. Note that our approach uses a pure software solution in an attempt to explore a technique that would work even in conjunction with larger volume, for which the use of texture mapping hardware may be less attractive. On the other hand, there is also a sizable class of volumes with sparse, yet space-filling features, such as arteries or nerve cells, where software rendering speed can actually match that of brute-force texture mapping hardware implementations[18].

Our paper is structured as follows. First, in Section 2, we provide some preliminary background and the theoretical aspects of our work, while Section 3 describes the actual implementation. Section 4 reports on the results, and Section 5 concludes with final remarks and pointers for future work.

2. Theory

2.1. Preliminaries

The low-albedo volume rendering integral can be written as follows:

$$r(u, v) = \int_0^L c(d(s))\tau(d(s))\exp\left(-\int_0^s \tau(d(t))dt\right) ds \quad (1)$$

where $r(u, v)$ is the value of the ray spawned at image coordinate (u, v) , $d(s)$ is the (interpolated) volume density at location s along the ray, and c and τ are the color and extinction mapped to d , respectively, via the transfer functions. Note, that c is a 3-vector of (r, g, b) . Since the continuous integral is, in the general case, not solvable analytically, it is commonly written and evaluated in discrete form:

$$r(u, v) = \sum_{i=0}^{L/\Delta s} c(d(i\Delta s))\tau(d(i\Delta s))\exp\left(-\sum_{j=0}^{i-1} \tau(d(j\Delta s))\Delta s\right)\Delta s \quad (2)$$

where Δs is the ray sampling interval. If we replace the exponential function by its Taylor series, we get the familiar compositing equation [13]:

$$r(u, v) = \sum_{i=0}^{L/\Delta s} c(d(i\Delta s))\alpha(d(i\Delta s))\left(\prod_{j=0}^{i-1} (1 - \alpha(d(j\Delta s)))\right) \quad (3)$$

where α is the sample opacity, normalized for sample distance when $\Delta s \neq 1$ [14].

Let us first concentrate on the continuous form of the low-albedo volume rendering integral given in (1), and the fact that we will be re-using the $d(s)$, i.e., the interpolated samples along the ray. This seems to be a simpler task than having to solve (1) directly from the raw volume, involving interpolation. Can we now obtain an approximate analytical solution for $r(u, v)$, or at least a fast discrete one? We shall investigate this next.

An effective way to go about the problem is to try and decompose c , τ (and perhaps also d) into sets of basis functions, and then just modify the weights of these basis functions according to a new transfer function setting before recombination. One such decomposition is the Fourier transform, or the related cosine transform. Kaneda et.al. [9] attempted this a few years ago, but only succeeded for the color transfer function. For the Fourier approach to work, the inner integral of (1) needs to remain fixed as a multiplicative constant, and hence the opacity transfer function remains fixed as well, which prevents the Fourier method from being able to change the transparencies of exterior and interior structures on the fly.

But the Fourier series is not the only mechanism to obtain a decomposition into basis functions. Another set of basis functions can be obtained by a series of polynomials, which is described next.

2.2. Polynomial fit

We may express the transfer functions as well as the list of sample densities along a ray as polynomial functions, which could be obtained by a least squares fit. We will get:

$$\begin{aligned} c(d) &= a_n d^n + a_{n-1} d^{n-1} + \dots + a_0 = p_1^n \\ \tau(d) &= b_n \tau^n + b_{n-1} \tau^{n-1} + \dots + b_0 = p_2^n \\ d(s) &= e_n s^n + e_{n-1} s^{n-1} + \dots + e_0 = p_3^n \end{aligned} \quad (4)$$

Recall that c is a 3-vector, and so will be the a -coefficients. Note that we will only need to find the coefficients for the $c(d)$ and $\tau(d)$ on the fly, the coefficients for $d(s)$, on the other hand, will be computed in a pre-processing step. We shall now express (1) with these polynomials:

$$r(u, v) = \int_0^L p_1^n(p_3^n(s)) \cdot p_2^n(p_3^n(s)) \exp\left(-\int_0^s p_2^n(p_3^n(t))dt\right) ds \quad (5)$$

$$r(u, v) = \int_0^L p_4^{2n}(s) \cdot p_5^{2n}(s) \exp\left(-\int_0^s p_6^{2n}(t)dt\right) ds \quad (6)$$

$$r(u, v) = \int_0^L p_7^{4n}(s) \exp(-p_8^{2n+1}(s)) ds \quad (7)$$

Using the first two terms of the Taylor series expansion of $\exp(x)$: