

# Translucent and Opaque Direct Volume Rendering for Virtual Endoscopy Applications

Michael Meißner and Dirk Bartz

WSI/GRIS, University of Tübingen  
Auf der Morgenstelle 10/C9  
D72076 Tübingen, Germany  
Email: {meissner,bartz}@gris.uni-tuebingen.de

**Abstract.** Virtual endoscopy applications frequently require the visual representation of several material interfaces to show the relevant data feature to the user. This requires the specification of complex transfer function which classify the various materials and color them appropriately.

In this paper, we explore the use of the direct volume rendering for virtual endoscopy. We specifically look into the visual representation of different anatomical features of various volume datasets, which are located below the inner surface of the organ of interest. Furthermore, we present how interactivity can be accomplished with the VIZARD II ray casting accelerator board.

**Keywords:** Direct Volume Rendering, Indirect Volume Rendering, Virtual Endoscopy.

## 1 Introduction

Virtual endoscopy is one of the most active fields of medical applications in computer graphics. Several different rendering and navigation techniques are used for the virtual examination of a variety of body organs.

Standard graphics hardware is used to render polygonal surface models [27, 18, 11, 3], extracted with the Marching Cubes algorithm [17]. In contrast, volume-rendering techniques are used, partially for better visual quality, partially for interactive speed [24, 31, 8, 1]. Unfortunately, interactive speed was always compromising visual quality, general applicability, or flexibility. In [24] and [5], key-framed animations are generated offline, which frequently leads to the time-intensive refinement of the key-framed animation. You et al. used a 16 processor SGI Challenge for parallel volume-rendering of isosurfaces [31]. In contrast, Gobetti et al. used the 3D texture mapping hardware abilities of high-end graphics systems for volume rendering. However, the lack of shading reduced the visual quality significantly [8]<sup>1</sup>. Furthermore, the size of the texture memory limits the size of datasets severely, while swapping techniques like

<sup>1</sup> In 1998 and 1999, several approaches were presented which enable isosurface shading [30] and volumetric shading [19] using 3D texture mapping. However, this approach does not provide sufficient performance for interactive endoscopy applications.

bricking reduce the framerate. The Navigator software of General Electric uses isosurface ray casting with approximately one frame per second. Even if the performance of the 1996 results has significantly improved, it hardly can be viewed as interactive [6]. Similarly, the Siemens Medical Systems Virtuoso workstation uses a 2D texture mapping approach for volume rendering [10]. However, this technique provides only reduced image quality – due to bilinear interpolation and the lack of gradient-based shading – at a low framerate. The Virtuoso workstation now also provides an option of using RTViz’s VolumePro board [23]. Unfortunately, VolumePro does not support perspective projection, which is absolutely mandatory for endoscopic explorations, and it will not be included into the functionality of VolumePro in the near future. A voxel-slab based multi-pass approach to simulate a perspective projections for VolumePro has been proposed by Wan et al. [29]. While no image quality or render performances data is presented in the web version of that paper, Li et al. [16] speak of five fps at low image quality for a  $256^3$  voxel volume. In their own paper, Li and Kaufman describe another voxel-slab based system which uses image-based rendering techniques (image warping) to reduce artifacts and to provide a high framerate [16].

In 1998, VIZARD II was introduced, a flexible architecture which implements a true ray casting approach with perspective and parallel projections [21]. While VIZARD II is currently only available as simulation, its first prototype of a hardware implementation is scheduled for completion in the second quarter of 2001. In this paper, we present the results of the virtual endoscopy system VIVENDI [3] which uses the VIZARD II simulation to render virtual flythroughs to various organs from “real life” patient datasets. The rendering performance will be estimated using the VIZARD II simulator, which allows a cycle accurate evaluation.

Aliasing problems due to undersampled areas in perspective projections were already addressed by a number of people. Novins et al. [22] proposed a ray casting scheme which splits a ray in two, once the ray divergence causes undersampling. A similar approach was later presented by Kreeger et al. [13]. A different approach was followed by Levoy and Whitaker [15] and Swan et al. [25]; they used a mipmap-like representation to account for diverging rays [15] or insufficient support of the sampling kernels for splatting [25]. Other approaches suggest pre-integrated volume rendering to reduce sampling artifacts [7]. Here, we focus on the examination of the effects of oversampling and filtering, in the context of VIZARD II.

Our paper is organized as follows; in the next section, we briefly outline the functionality and performance of VIZARD II. In Section 3, we present the visual results of the experiments. In particular, we focus on the influence of sampling and filtering due to the transfer functions of opaque and transparent volume rendering. Finally, we summarize our paper in Section 4.

## **2 VIZARD II**

The VIZARD II is a special purpose PCI card consisting of several components. The main design goal has been the implementation of a ray casting algorithm as well as a

high degree of flexibility to enable future changes and other implementations. In the following, the underlying algorithm, the architecture, and the implemented units are described.

## 2.1 Volume Rendering Algorithm

The algorithm implemented on the VIZARD II system mainly follows the work presented in [14], implementing a full ray casting pipeline. Rays are cast through a volumetric, possibly non-uniform, regular dataset. To ensure high image quality, sampling needs to be freely selectable in each dimension. Different rendering modes such as MIP, unshaded, shaded, etc. are supported.

Starting with the viewing parameters such as eye position, view direction, view up vector, etc., the position of the view plane is calculated. For each point  $P_{i,j}$  of the view plane, a ray is cast into the volume and tested if it hits the volume data. A sample is generated by trilinearly interpolating the eight neighboring voxels on the grid. In a similar fashion, the gradient at sample location is computed. Instead of computing gradients at voxel location on the fly — which would result in a 32 gradient neighborhood per sample for a central difference gradient operator — gradients are considered to be a voxel property. This is similar to surface rendering where a normal is a vertex property and not computed from neighboring triangles on the fly. Another reason for this is that numerous gradient operators exist but each of them would require a different memory interface to deliver the required data.

Generally, memory access is a crucial aspect in all volume rendering architectures. Derived from [12] and similar to VIRIM [9] and others, an eight-way interleaved memory is used for the VIZARD II system. In contrast to previous approaches, SDRAM DIMM memory modules are used to allow different volume memory sizes without the need of fabricating a new PCI card. Since DIMMs come in modules providing a 64 bit data bus (72 bit including eight parity bits), four DIMM modules are used spending 32 bit per voxel and replicating volume data in one dimension to exploit the remaining 32 bits.

Classification is performed using the sample value for the look-up into the classification table which is realized as two 32 bit SRAMs. The result of addressing the two SRAMs is a  $(r, g, b, \alpha, k_a, k_d, k_s)$  tuple. Phong shading is performed using the gradient at sample position and a look-up table based shading technique as presented in [28, 26]. The tables<sup>2</sup> require 3 KBytes of memory and only need to be computed when the illumination parameters change, i.e. the direction of the lights. The obtained diffuse and specular shading intensities are multiplied with the material properties and with the color of the classified sample.

As a last step, each classified and shaded sample needs to be composed with the previously accumulated color. The final pixel of a ray is obtained once the last pixel of the ray is composed or once the accumulated opacity is higher than a certain threshold, i.e.

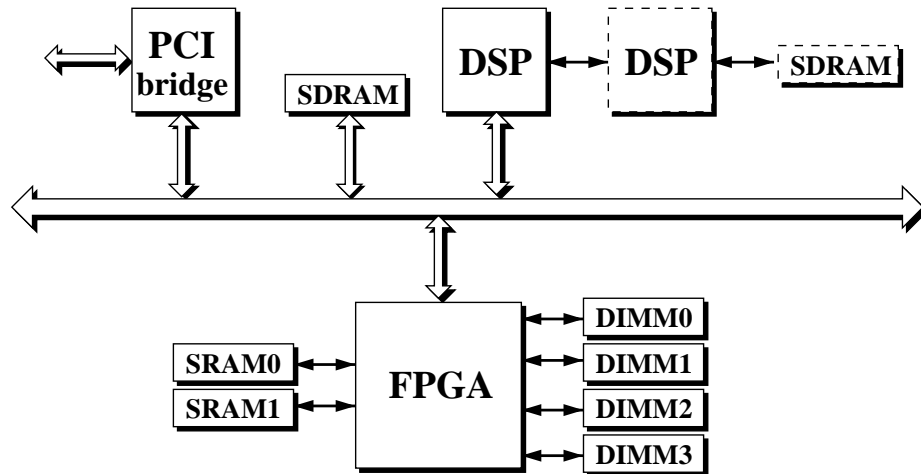
---

<sup>2</sup> One for the diffuse and one for the specular intensity which can be eye point independent.

$\alpha \geq 0.98$ . For the presented views of endoscopic applications, early ray termination is a powerful acceleration technique increasing the overall framerate significantly. However, when classifying the wall of ventricle as semi-transparent to visualize the arteries, early ray termination performs less efficient.

## 2.2 Architecture

A schematic overview of the VIZARD II system architecture is given in Figure 1. The



**Fig. 1.** VIZARD II system architecture: The second DSP (dashed block) is optional.

VIZARD II system architecture has been designed for a ray casting algorithm and therefore has a dedicated memory interface to provide optimal voxel access for arbitrary rays. A local bus is used to transfer data within the system, but also to enable data transfer to and from the outside world (PCI bridge). The main component is the reconfigurable FPGA chip. It controls two SRAM and four DIMM modules. Furthermore, there is one DSP and a SDRAM which is the external memory of the DSP. A second DSP and SDRAM are optional and not needed for the implementation of ray casting<sup>3</sup>.

## 2.3 Performance

The PCI bus running at 33 MHz is only heavily used during downloading the volume data onto the card. While static datasets or fixed sequences of volumes can be stored

<sup>3</sup> The VIZARD II board has been designed in a joint project with Phillips Research Hamburg including an implementation of a volume reconstruction algorithm which makes use of the second DSP.

initially on the card, real time volume updates of entire volumes is currently not possible.

Generally, data is sent to the board and stored in the DIMM modules. The transfer of a dataset of  $256^3$  voxels over the PCI bus takes roughly 0.15 seconds. The classification tables are also sent over the PCI bus but are neglectable since they are only 2 KBytes in size. Even if the classification changes for every frame, 30 frames per second would require 60 KByte/s bandwidth which can easily be handled. Transferring an entire image back to the host requires 65 KByte/s which results in almost 8 MByte/s for 30 frames per second. This can also be handled by the PCI bus and does not introduce any bottle-neck.

The overall performance limitation of the system is given by its memory interface. The used DIMM modules run at 100 MHz and need 70 nsec for a pre-charge and row activate. In average, this results in 12.7 nsec for each sample. Thus, 80 million trilinearly interpolated samples can be generated per second. For an one to one mapping of samples to voxels and a dataset of  $256^3$  voxels, this translates into 5 frames/s. However, for flythroughs, as presented in this paper, the framerate is much higher. Additionally, dependent on the classification the overall framerate can be increased significantly due to early ray termination. Generally, early ray termination is a very powerful technique in endoscopic applications. For the presented views where the ventricle is classified opaque, framerates well above 20 can be accomplished which is certainly enough for endoscopic applications. Please note, even though the VolumePro system of Mitsubishi can deliver sustained 30 frames per second, correct perspective projections mandatory for endoscopic applications are not available. This is due to the architecture of VolumePro which requires that rays are sent parallel to each other following an one to one mapping of samples to voxels. As already mentioned in Section 1, voxel-slab based multi-pass approaches only simulate a perspective projections at significant lower framerates and low image quality [29].


### 3 Experiments

In this section, we present an MRI datasets which contains multiple data features. The dataset is generated by an MRI TOF (Time of Flight) sequence which represents the cerebrospinal fluid (CSF) filled ventricular system with a very low intensity. In contrast, the MRI TOF sequence emphasizes moving particles which are not yet saturated by the magnetic resonance. This results in a high intensity of voxels associated with blood vessels. Figure 4 shows the opaque rendered material interface between the low intensity third ventricle and the surrounding brain tissue. In Figures 5 and 6, we substantially reduced the opacity value of the this material interface and introduced a full opaque red material, which represents the blood vessels in viewing direction.

The additional step in the opacity and color transfer functions for the translucent rendering for Figure 6 introduces additional higher frequencies compared to the transfer functions for the (mostly) opaque rendering in Figure 4. These higher frequencies

require a higher sampling rate to reduce the aliasing artifacts, i.e., staircase artifact, although some of the artifacts are already reduced due to low transparency. This problem aggravates due to the close position of the view point to the material interface.

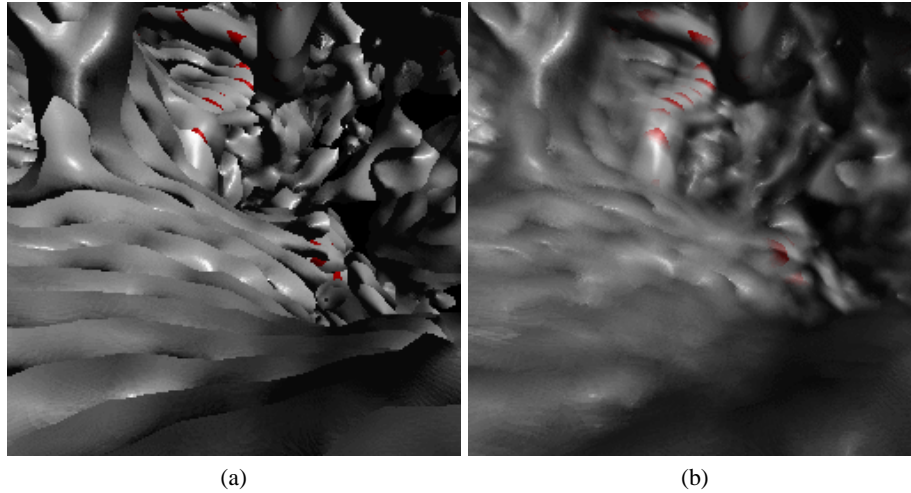
These sampling problems can be addressed in many ways (see Section 1). In this paper, we basically explore two different ways. The first obvious solution is to increase the rate of samples along the rays through the volume. Depending on the distance to the material interface in perspective views, oversampling in  $z$  (along the rays) of more than ten samples per unit distance can be required. Figure 2 shows two images rendered at uni-sampling (a) and eight-times oversampling (b), exhibiting the image quality differences.



**Fig. 2.** Influence of sampling rate of images rendered with binary opacity transfer functions; (a) uni-sampling, (b) eight-times oversampling.

The second possible solution is a modification of the transfer functions, in particular of the opacity function. An oversampling of up to ten is only required if binary opacity changes are specified. If linear, quadratic, or even higher order (in contrast to zero-order binary opacity functions) opacity functions are used, the oversampling rate can be reduced. However, higher order opacity classification also introduces blurring [20], due to its filter properties, if its support is too large. Figure 3 shows two images which use a binary (a) and a ramp (b) opacity transfer function.

Overall, a combination of oversampling along the rays and higher-order classification (in contrast to zero-order binary opacity transfer functions) with a limited filter width is used (see Fig. 4 - 6). A higher oversampling rate also introduces higher rendering costs of VIZARD II. If we assume four-times oversampling, VIZARD II achieves an estimated framerate of five fps for opaque classifications. If we choose a more costly



**Fig. 3.** Influence of binary (a) and ramp opacity transfer functions (b) with uni-sampled volumes.

transparent rendering (more samples), the implicit averaging step can be traded with a further reduction of the sampling rate, thus achieving a similar framerate. For a better visualization of the surface of the material interface between the ventricle wall and the brain tissue, we can limit the number of contributing samples, once the material interface is detected. We expect images similar to Figure 5, where we rendered two iso-surfaces and blended them subsequently with appropriate opacities. As another positive effect, this techniques would also increase the rendering performance.

### 3.1 Voxels versus Polygons

Alternative polygonal representations of the data can also be used with a standard graphics accelerator. However, isosurface extraction algorithms like Marching Cubes [17] tend to generate many, small triangles which pose a significant rendering load a graphics accelerator. In [4], we modified the VIVENDI system to address this kind of data. Due to segmentation problems of the noisy MRI TOF data, we needed to manually register and match two different MRI sequences for the ventricular system (MRI TSE), and for the arterial blood vessels (MRI TOF). Direct volume rendering provides additional filter steps of the transfer functions which reduce the noise in the data. It also provides volumetric cues which indicate the volumetric distance between the boundary of the ventricles, and the blood vessels. For a more detailed discussion of the visual differences we refer the interested reader to [2].

## 4 Conclusion and Future Work

In this paper, we explored the use of ray casting for opaque and translucent rendering of segmented volume datasets for virtual endoscopy applications. We noticed that oversampling is in particular mandatory to accomplish good image quality, at least for segmented datasets as used in this case. Due to the segmentation, the walls of the ventricle are very thin and require a high sampling frequency. A sampling distance of one yields to insufficient image quality due to undersampling. In general, we found that a unit sampling distance of 0.25 is necessary to accomplish satisfactory image quality. For datasets which do not have such high frequencies, the Nyquist frequency suffices to accomplish good image quality. The higher sampling rate can also be reduced by using higher order classification functions, since this introduces an additional filter operation.

VIZARD II is a special purpose hardware accelerator for true ray casting and high image quality. With its highly optimized memory interface, it is capable of generating up to 80 million trilinearly interpolated samples and gradients per second using one processing pipeline only. Each sample is Phong shaded using per sample material properties and composed at high precision to ensure highest image quality. With respect to endoscopic applications, this performance can deliver about 20 frames per second (fps) for uni-sampled opaque rendering of the ventricle since early ray termination can be exploited. With four-times oversampling, this framerate is reduced to estimated five fps.

Future work will focus on how to avoid such extremely high sampling frequencies in empty space by applying space leaping. Furthermore, we would like to investigate how to avoid slicing patterns by jittering the sampling position from ray to ray.

### Acknowledgments

This work has been supported by the Workstations Systems Lab of Hewlett Packard, Ft. Collins, by project 382 and CatTrain of the German Research Council (DFG).


The MRI head dataset was provided by the Department of Neuroradiology of the University Hospital Tübingen.

### References


1. D. Auer and L. Auer. Virtual Endoscopy - A New Tool for Teaching and Training in Neuroimaging. *International Journal of Neuroradiology*, 4:3–14, 1998.
2. D. Bartz and M. Meißner. Voxels versus Polygons: A Comparative Approach for Volume Graphics. In *Proc. of Volume Graphics*, pages 33–48, 1999.
3. D. Bartz and M. Skalej. VIVENDI - A Virtual Ventricle Endoscopy System for Virtual Medicine. In *Proc. of Symposium on Visualization*, pages 155–166, 324, 1999.
4. D. Bartz, W. Straßer, Ö. Gürvit, , D. Freudenstein, and M. Skalej. Interactive and Multimodal Visualization for Neuroendoscopic Interventions. In *Proc. of Symposium on Visualization*, 2001.

5. J. Beier, T. Diebold, H. Vehse, G. Biamino, E. Fleck, and R. Felix. Virtual Endoscopy in the Assessment of Implanted Aortic Stents. In *Proc. of Computer Assisted Radiology*, pages 183–188, 1997.
6. C. Davis, M. Ladds, B. Romanowski, S. Wildermuth, J. Knoploch, and J. Debatin. Human Aorta: Preliminary Results with Virtual Endoscopy Based on Three-dimensional MR Imaging Data Sets. *Radiology*, 199:37–40, 1996.
7. K. Engel, M. Kraus, and T. Ertl. High-Quality Pre-Integrated Volume Rendering Using Hardware-Accelerated Pixel Shading. In *Proc. of Eurographics/SIGGRAPH Workshop on Graphics Hardware*, 2001.
8. E. Gobbetti, P. Pili, A. Zorcolo, and M. Tuveri. Interactive Virtual Angioscopy. In *Proc. of IEEE Visualization*, pages 435–438, 1998.
9. T. Guenther, C. Poliwoda, C. Reinhard, J. Hesser, R. Männer, H. Meinzer, and H. Baur. VIRIM: A Massively Parallel Processor for Real-Time Volume Visualization in Medicine. In *Proc. of Eurographics Workshop on Graphics Hardware*, pages 103–108, 1994.
10. Ö. Gürvit, M. Skalej, R. Riekmann, U. Ernemann, and K. Voigt. Rotational Angiography and 3D Reconstruction in Neuroradiology. *electro medica*, 68(1):31–37, 2000.
11. L. Hong, S. Muraki, A. Kaufman, D. Bartz, and T. He. Virtual Voyage: Interactive Navigation in the Human Colon. In *Proc. of ACM SIGGRAPH*, pages 27–34, 1997.
12. G. Knittel. Verve - Voxel Engine for Real-time Visualization and Examination. *Computer Graphics Forum*, pages 37–48, sep 1993.
13. K. Kreeger, I. Bitter, F. Dachille, and A. Kaufman. Adaptive Perspective Ray Casting. In *Proc. of Symposium on Volume Visualization*, pages 55–62, 1998.
14. M. Levoy. *Display of Surfaces From Volume Data*. PhD thesis, Department of Computer Science, University of North Carolina at Chapel Hill, 1989.
15. M. Levoy and R. Whitaker. Gaze-Directed Volume Rendering. In *Symposium on Interactive 3D Graphics*, pages 217–223, 1990.
16. W. Li and A. Kaufman. Real-Time Volume Rendering for Virtual Colonoscopy. In *Proc. of Volume Graphics*, 2001.
17. W. Lorensen and H. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In *Proc. of ACM SIGGRAPH*, pages 163–169, 1987.
18. W. Lorensen, F. Jolesz, and R. Kikinis. The Exploration of Cross-Sectional Data with a Virtual Endoscope. In R. Satava and K. Morgan, editors, *Interactive Technology and New Medical Paradigms for Health Care*, pages 221–230. IOS Press, 1995.
19. M. Meißner, U. Hoffman, and W. Straßer. Enabling Classification and Shading for 3D Texture Mapping Based Volume Rendering. In *Proc. of IEEE Visualization*, pages 207–214, 1999.
20. M. Meißner, J. Huang, D. Bartz, K. Mueller, and R. Crawfis. A Practical Evaluation of Four Popular Volume Rendering Algorithms. In *Proc. of Symposium on Volume Visualization and Graphics*, pages 81–90, 2000.
21. M. Meißner, U. Kanus, and W. Straßer. VIZARD II, A PCI-Card for Real-Time Volume Rendering. In *Proc. of Eurographics/SIGGRAPH Workshop on Graphics Hardware*, pages 61–68, 1998.
22. K. Novins, F. Sillion, and D. Greenberg. An Efficient Method for Volume Rendering using Perspective Projection. In *Workshop on Volume Visualization*, pages 95–102, 1990.
23. H. Pfister, J. Hardenbergh, J. Knittel, H. Lauer, and L. Seiler. The VolumePro Real-Time Ray-Casting System. In *Proc. of ACM SIGGRAPH*, pages 251–260, 1999.
24. R. Shadidi, V. Argiro, S. Napel, L. Gray, H. McAdams, G. Rubin, C. Beaulieu, R. Jeffrey, and A. Johnson. Assessment of Several Virtual Endoscopy Techniques Using Computed Tomography and Perspective Volume Rendering. In *Visualization in Biomedical Computing*, volume LNCS 1131, pages 521–528, 1996.


25. E. Swan, K. Müller, T. Möller, N. Shareef, R. Crawfis, and . Yagel. An Anti-Aliasing Technique for Splatting. In *Proc. of IEEE Visualization*, pages 197–204, 1997.
26. J. Terwisscha van Scheltinga, J. Smit, and M. Bosma. Design of an on Chip Reflectance Map. In *Proc. of Eurographics Workshop on Graphics Hardware*, pages 51–55, 1995.
27. D. Vining, D. Stelts, D. Ahn, P. Hemler, Y. Ge, G. Hunt, C. Siege, D. McCorquodale, M. Sarojak, and G. Ferretti. FreeFlight: A Virtual Endoscopy System. In *First Joint Conference, Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, volume LNCS 1205, pages 413–416, 1997.
28. D. Voorhies and J. Foran. Reflection Vector Shading Hardware. In *Proc. of ACM SIGGRAPH*, pages 163–166, 1994.
29. M. Wan, W. Li, K. Kreeger, I. Bitter, B. Chen, A. Kaufman, Z. Lian, D. Chen, and M. Wax. 3D Virtual Colonoscopy with Real-time Volume Rendering. In *SPIE Symposium on Medical Imaging*, 2000.
30. R. Westermann and T. Ertl. Efficiently Using Graphics Hardware in Volume Rendering Applications. In *Proc. of ACM SIGGRAPH*, pages 169–177, 1998.
31. S. You, L. Hong, M. Wan, K. Junyaprasert, A. Kaufman, S. Muraki, Y. Zhou, M. Wax, and Z. Liang. Interactive Volume Rendering for Virtual Colonoscopy. In *Proc. of IEEE Visualization*, pages 343–346, 1997.



**Fig. 4.** Snapshots of the opaque third ventricle from a ventricular system flythrough at varying sampling rates in  $z$  (along the rays); The three different rows show different sampling rates; (I) uni-sampling, (II) two-times oversampling, (III) four-times oversampling. The two columns show images rendered with binary (left column) and ramp opacity transfer functions (right column).



**Fig. 5.** (a) Snapshot of the transparent third ventricle with opaque red arterial circle of Willis with a ramp opacity transfer functions. The individually opaque rendered images are blended by image-processing techniques (without absorption) with 40% and 60% opacity. (b) Transfer functions for the surface of the third ventricle, and for the surface of the blood vessels (c).



**Fig. 6.** (I) Snapshots of the transparent third ventricle with opaque red arterial circle of Willis with a ramp opacity transfer function. The images are composited using regular direct volume rendering, taking absorption into account; we use four-times oversampling. (II) Transfer functions of the respective images. The left column shows a more solid surface of the third ventricle than the right column.