Volume Rendering

Marco Ament

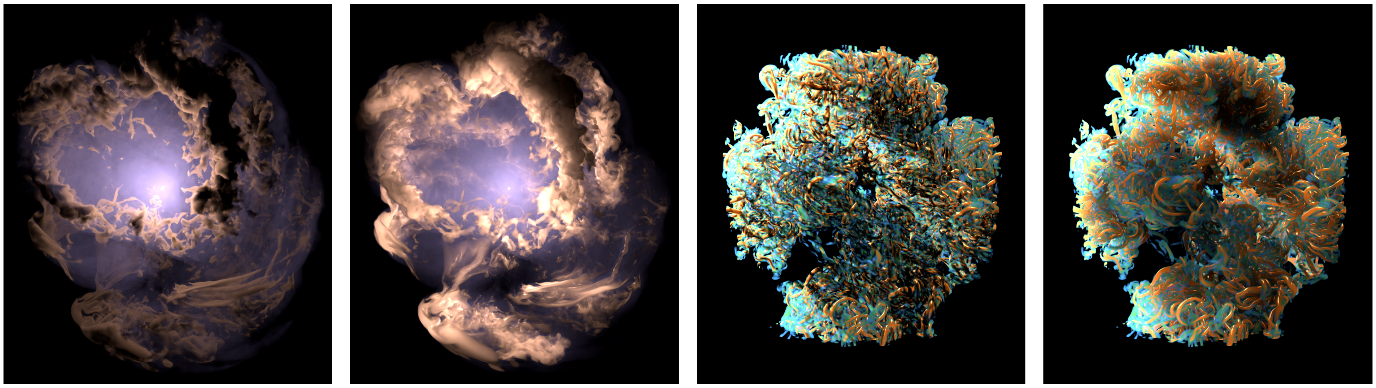

(a) (b) (c) (d)

Fig. 1: (a),(b) Volume renderings of a supernova simulation with different optical models. (a) Single scattering. (b) Ambient volume scattering [7]. (c),(d) Volume renderings of a flow simulation with different optical models. (c) Single scattering. (d) Low-pass filtered volumetric shadows [5].

Abstract—This report summarizes my research on the visualization of scalar fields and the rendering of participating media. A strong focus of my work is on the development of optical models for exploring and visualizing volumetric data sets interactively. In particular, my research is concentrated on goal-oriented optical models that offer simple control over the visual properties of a visualization combined with high performance. In addition, the perception of spatial depth and size of volumetric, semi-transparent features plays an important role in my work. For this purpose, I study and develop new methods to efficiently illuminate volumetric data sets with a focus on intuitive parameter control. With physically-based rendering of participating media, the optical models are strictly given by the laws of radiative transfer and high fidelity is a major goal. In this field, my work introduces an optical model from physics to graphics that rigorously describes the complex light transport in participating media with a spatially varying index of refraction. For visualizing large data sets on distributed clusters, I also develop parallel volume rendering algorithms, focusing on load balancing and caching strategies.Finally, I also work on parallel algorithms that reconstruct a volumetric model from a single input image to visualize astronomical objects like supernovae or emission nebulae with volume rendering.

Index Terms—Direct volume rendering, direct interval volume visualization, sort first parallel volume rendering, ambient volume scattering, low-pass filtered volumetric shadows, anisotropic ambient volume shading, extinction-optimized volume illumination, ambient volume illumination, refractive radiative transfer equation, reconstruction of astronomical objects

1 Introduction

Scalar fields play a fundamental role in many scientific disciplines and applications. The increasing computational power offers scientists and digital artists novel opportunities for complex simulations, measurements, and models that generate large amounts of data. Visualization is an essential interface between the usually abstract numerical data and human operators who want to gain insight [1, 4].

Depending on the application focus, the different requirements on a visualization or rendering must be considered in the development of novel techniques. In scientific volume visualization, spatial perception and interactive controllability of the visual representation are usually more important than physical accuracy. For example, in Section 2, a novel model is presented for mapping opacity to isosurfaces that have a small but finite extent. Compared to physically based opacity, the presented approach offers improved control over occlusion and visibility of such interval volumes.

• *Marco Ament is with Karlsruhe Institute of Technology, Germany. E-mail: ament@kit.edu.*

Moreover, my research is focused on the development of illumination models for interactive volume rendering. In Section 3, an algorithm for the approximation of multiple scattering is presented to achieve advanced lighting effects like soft shadows and translucency. The main benefit of this contribution is an improved perception of volumetric features with full interactivity of all relevant parameters. Section 4 presents an efficient algorithm for low-pass filtering of volumetric shadows to reduce disruptive hard shadow edges and noise-like visual patterns in the presence of fine structures due to single scattering illumination. In Section 5, a novel approach is introduced to illuminate anisotropic structures in volumetric data sets. Typically, isotropic illumination models like Blinn–Phong are used for surface shading. However, anisotropic shading models can be beneficial to visually convey the shape of anisotropic structures like edges. Section 6 presents a discussion on different illumination algorithms that are based on the ambient region of each point in the data set to achieve improved perception of spatial depth compared to pure local illumination models. A common problem of many illumination methods in volume visualization is to transmit enough light to important features in the data set. In Section 7, an optimization algorithm is presented that allows one to define important features and then to illuminate them properly. At the same time, these important structures become also more visible from the camera.

In visual media, scalar fields often describe complex materials and their realistic appearance is of highest interest by means of accurate rendering models and algorithms. In this case, an accurate description of light transport in the real world is essential for realistic image synthesis of natural phenomena. In particular, physically based rendering aims to produce predictive results for real material parameters. Section 8 presents a physically based light transport equation for inhomogeneous participating media that exhibit a spatially varying index of refraction. In addition, an extended photon mapping algorithm is introduced that provides a solution of this optical model.

In addition to the visual representation, the continuously growing data set sizes pose challenges with respect to performance and data scalability. In particular, fast graphics processing units (GPUs) play a central role for current and future developments in distributed rendering and computing. For volume visualization, Section 9 presents a parallel algorithm that dynamically decomposes image space and distributes work load evenly among the nodes of a multi-GPU cluster. The presented technique facilitates illumination with volumetric shadows and achieves data scalability with respect to the combined GPU memory in the cluster domain.

The 3D visualization of astronomical nebulae is an important application of volume rendering. However, volumetric models of these nebulae are difficult to obtain and requires significant modeling effort. Another strategy is to reconstruct a 3D model automatically from telescope images, but since only one perspective is available, it is challenging to obtain plausible results. In Section 10, a novel algorithm is presented that reconstructs plausible models of emission nebulae which exhibit approximate symmetries. With the help of multi-GPU cluster, even high-resolution models can be computed within a few hours. The volumetric data sets can then be rendered interactively with direct volume rendering.

2 Direct Interval Volume Visualization

The paper *Direct Interval Volume Visualization* [9] introduces a novel optical model, based on emission and absorption, for visualizing interval volumes and isosurfaces with scale-invariant opacity. The introduced technique is derived from the visual properties of isosurfaces, which means that opacity does not vary with the viewpoint and that scale has no influence. In standard direct volume rendering (DVR), opacity is strongly affected by the distance that a ray traverses in the spatial and data domain. With the introduced model, this potentially strong imbalance of opacity is avoided by developing an algorithm that compensates varying scale in both domains. In this way, the visual properties of interval volumes are closer to the visual properties of isosurfaces than of large volumetric regions. It is shown that the novel model for finite interval volumes is also able to visualize true isosurfaces with constant opacity, which effectively generalizes previous approaches. Furthermore, since the quality of the visualization strongly depends on accurate sampling, a novel and efficient method is introduced to analytically reconstruct the scalar field inside a cubic cell of a uniform grid. By assuming trilinear interpolation, a cubic polynomial is locally computed from only four samples in each cell by exploiting the fast texture unit of modern GPUs. In this way, crack-free rendering can be guaranteed in combination with higher frame rates than with previous reconstruction methods. This paper was published at IEEE SciVis 2010 [9].

In Figure 2, the orange isosurface is modeled with a thin but finite interval volume using the novel approach in both images. Furthermore, soft tissue (green), brain (red), and dentin (blue) are visualized with standard DVR. However, the cranium is rendered differently in both images. In Figure 2(a), standard DVR is employed and the brain area is almost transparent whereas the silhouette of the cranium is opaque. Furthermore, the tooth enamel is not visible at all. In contrast, in Figure 2(b), the cranium is modeled as an interval volume. The finite extent of the interval volume is apparent at the semi-transparent silhouette. The surface-like structures of the cranium are visible and the enamel is classified properly.

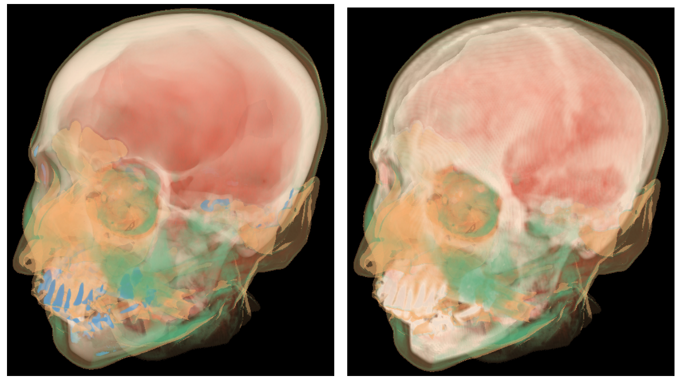

(a) (b)

Fig. 2: Volume renderings of a computer tomography of a human head with different optical models for the cranium. (a) Standard DVR (b) interval volume visualization.

3 Ambient Volume Scattering

The paper *Ambient Volume Scattering* [7] introduces a simplified optical model for the efficient illumination and interactive visualization of volumetric data sets. In participating media, far-range scattering effects often do not contribute much radiance to the final image due to the exponential attenuation. As a consequence, this paper presents a novel illumination model that confines multiple scattering into spheres of finite radii around any given sample point. A Monte-Carlo simulation preintegrates light transport for a set of material parameters like anisotropy and extinction. At rendering, a small preintegration table is accessed to efficiently estimate ambient light transport by a simple texture lookup. The presented approach effectively generalizes local shadowing techniques like volumetric ambient occlusion. In particular, the algorithm provides interactive volumetric scattering and soft shadows, with interactive control of the transfer function, anisotropy parameter of the phase function, lighting conditions, and viewpoint. Since preintegration does not depend on the data set or transfer function, the technique is also well suited for the visualization of time-dependent data. This technique was published in a conference paper at IEEE SciVis 2013 [7] and received an Honorable Mention award.

Figures 1(a) and (b) show volume visualizations of a supernova with different optical models. In Figure 1(a), a single point light source in the center illuminates the supernova with single scattering. The directional shadows provide important visual cues for the perception of spatial depth, but many features remain hidden in the dark regions. In Figure 1(b), the same point light source is employed with ambient volume scattering and gradient-based shading, which provides softer shadows and more interaction of light to obtain a well-lit visualization. The image in Figure 1(b) was selected as winner (1st place) of the Computer Graphics Forum Cover Image Contest 2014 [8].

4 Low-pass Filtered Volumetric Shadows

The paper *Low-pass Filtered Volumetric Shadows* [5] presents a novel and efficient method to compute volumetric soft shadows for interactive direct volume visualization to improve the perception of spatial depth. By direct control of the softness of volumetric shadows, disturbing visual patterns due to hard shadows can be avoided and users can adapt the illumination to their personal and application-specific requirements. Conceptually, the classified data set is convolved with a low-pass filter to compute volumetric soft shadows. To avoid expensive shadow ray marching, a three-dimensional summed area table is aligned with the light source, which enables efficient evaluation of soft shadows for each point in constant time without shadow ray marching and the softness of the shadows can be controlled interactively. The technique offers interactive control of the transfer function, light source positions, and viewpoint. The method was published in a conference paper at IEEE SciVis 2014 [5].

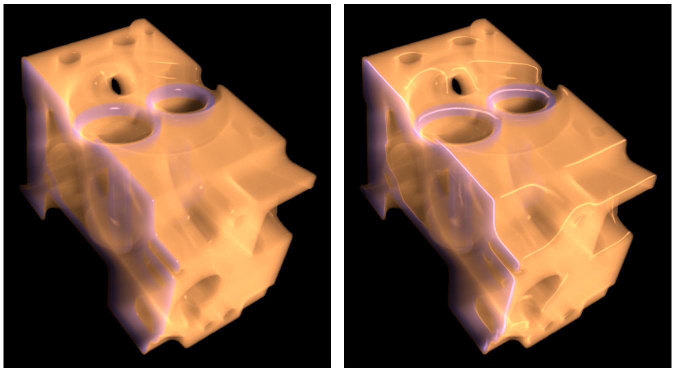

(a) (b)

Fig. 3: Visualizations of the Engine data set. (a) Isotropic shading with Blinn–Phong's BRDF. (b) Anisotropic shading with Lafortune's BRDF.

Figures 1(c) and (d) show volume visualizations of a flow simulation with different optical models. In Figure 1(c), two point light sources illuminate the data set with single scattering. The hard shadows and vortex features visually interfere with each other due to similar scale, hampering perception of depth. In Figure 1(d), a low-pass filter is employed, which significantly reduces the distracting streak patterns. The technique was also employed to render the image that won the 2nd place of the Computer Graphics Forum Cover Image Contest 2015 [6].

5 Anisotropic Ambient Volume Shading

The paper *Anisotropic Ambient Volume Shading* [3] introduces an algorithm to visualize the anisotropy of volumetric features with anisotropic bidirectional reflectance distribution functions (BRDFs). In volume visualization, local illumination is usually based on the isotropic Blinn–Phong model to compute the specular reflection of light on surface-like features. However, isotropic specular highlights do not support the perception of anisotropic features like elliptic or linear structures. This paper introduces an algorithm to automatically estimate the parameters for anisotropic BRDFs from the data set. In particular, the ambient region of each point in the data set is explored to measure the degree of anisotropy on a user-defined scale. Then, the anisotropy information is mapped to BRDF-specific parameters, which are used for local illumination at each shading point. In this way, the shape of the specular highlights visually encode the degree of anisotropy, supporting perception of shape and curvature. This technique was published in a conference paper at IEEE SciVis 2015 [3].

Figure 3 shows visualizations of the Engine data set. The data set exhibits many isotropic surfaces combined with sharp edges and curved concavities. The high transparency and the faint shadows make it difficult to visually convey salient structures of the engine block. In Figure 3(a), isotropic highlights are rendered with Blinn–Phong's BRDF. However, the highlights provide only few additional cues to help perceive the shape of the engine. In Figure 3(b), Lafortune's anisotropic BRDF is used to accentuate the gaskets or the edges of the engine block.

6 Ambient Volume Illumination

The paper *Ambient Volume Illumination* [10] discusses illumination algorithms for direct volume rendering that are based on the ambient region of each point in the data set. Ongoing research in volume visualization focuses on the development of simplified optical models that can reproduce the most important illumination effects that help improve perception, but on a phenomenological level to obtain interactive performance and simpler control of the parameters compared to accurate global illumination. With ambient occlusion, local shadows can be computed, which help perceive spatial depth, for example,

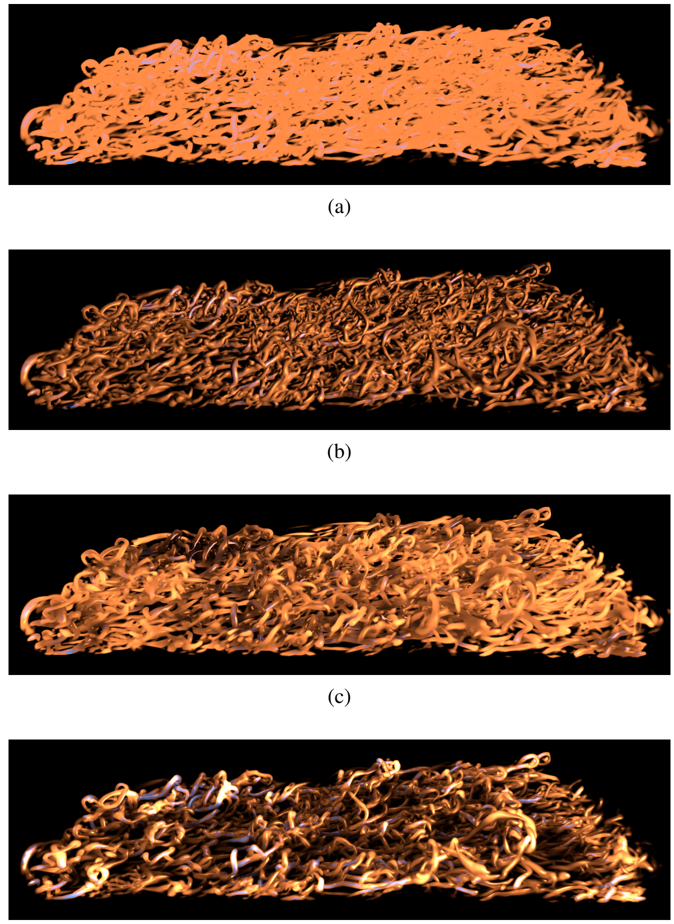

(a)

(b)

(c)

(d)

Fig. 4: Visualizations of a turbulent boundary layer using the λ_2 vortex criterion. (a) Emission–absorption model. (b) Blinn–Phong illumination model. (c) Ambient occlusion model. (d) Ambient scattering model.

to visually discover cavities in a larger volumetric structure. By going one step further, ambient scattering can be employed to simulate more complex interactions of light with matter, which can help perceive the shape of volumetric structures. Together, both algorithms provide means to compute ambient volume illumination with interactive performance and few, simple parameters. This work was published in a regular journal paper at IEEE Computing in Science and Engineering [10] in 2016.

Figure 4 shows visualizations of a simulation of a turbulent boundary layer with filament-like structures. With the emission–absorption model, in Figure 4(a), it is difficult to visually identify salient structures due to missing illumination. In Figure 4(b), Blinn–Phong's local illumination model is used, which helps perceive the individual tube structures, but perception of spatial depth is still limited. Even ambient occlusion, in Figure 4(c) has difficulties to visually convey the spatial arrangement of the features. In contrast, in Figure 4(d), ambient scattering provides more visual cues for spatial depth by means of directional shadows and indirect light.

7 Extinction-Optimized Volume Illumination

The paper *Extinction-Optimized Volume Illumination* [11] presents an algorithm to optimize the visibility and illumination of important features in volumetric data sets. In 3D volume visualization with single scattering illumination, important features are often occluded by surrounding structures. As a consequence, these features do no receive much light from external light sources and are hardly visible from the virtual camera. This paper introduces an optimization method to trans-

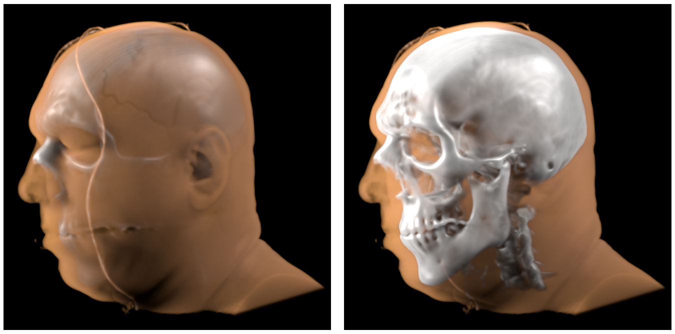

(a) (b)

Fig. 5: Visualizations of the Visible Male data set. (a) Standard single scattering. (b) The bone structures are mapped to high importance values and are visualized with optimized visibility and illumination.

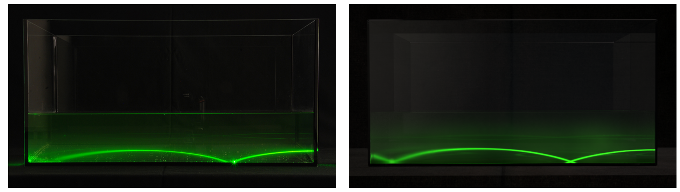

(a) (b)

Fig. 6: Bouncing laser beam due to continuous refraction in a glass tank filled with a sugar solution of varying concentration. (a) Photograph of real-world experiment. (b) Rendering of virtual experiment.

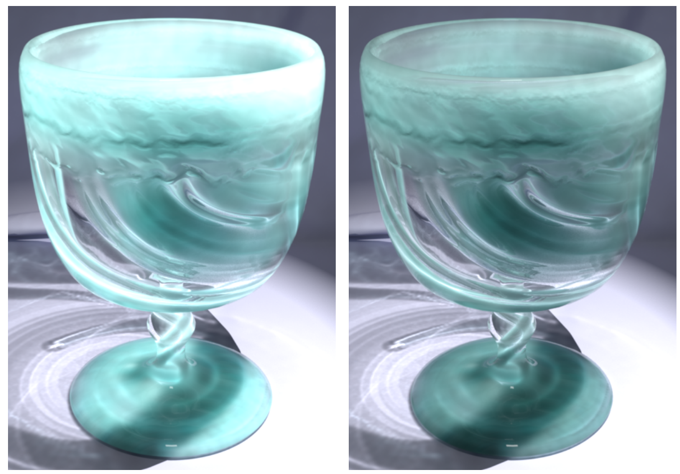

(a) (b)

Fig. 7: Renderings of a decorative cup made of a heterogeneous mix of light crown glass ($n = 1.48$) and jadeite ($n = 1.68$). (a) Wrong and (b) correct conservation of energy.

mit more light from and to important features than to less important structures. The algorithm is based on a user-defined function that controls the importance of any point in the data set. Together with the information from a regular transfer function, an optimization problem is formulated for each ray of the image synthesis process. Most important and in contrast to previous techniques, the presented optimization problem can be solved analytically, which is important to reach interactive performance. The algorithm is well-suited for efficient processing on the GPU and does not require any precomputation step. This work was published in a regular journal paper at IEEE Transactions on Visualization and Computer Graphics [11] in 2016.

Figure 5 shows visualizations of the Visible Male data set to demonstrate the benefit of optimizing both the visibility and illumination of important features. For comparison, Figure 5(a) is rendered with the standard single scattering model. In Figure 5(d), the bone structures are mapped to a higher importance value than the rest of the data set. As a consequence, the bone structures become visible and self-shadowing effects are still present, which provide important visual cues for perception of spatial depth.

8 Refractive Radiative Transfer Equation

The paper *Refractive Radiative Transfer Equation* [2] presents a generalization of the well-known radiative transfer equation to heterogeneous participating media that have a spatially varying index of refraction, represented as a continuous scalar field. Basic principles from non-linear Hamiltonian optics are reviewed to introduce a physically based light transport equation to the graphics and visualization community that accounts for curved trajectories and, in particular, for correct conservation of energy by means of an extended definition of radiance. A complete formal derivation of the optical model is presented for full reproducibility of all steps. As a secondary contribution, an extension of the photon mapping algorithm is presented that solves the novel model accurately. It is shown that previously published rendering approaches fail to conserve energy in such complex media. The more general light transport equation also facilitates the accurate rendering of time-dependent light transport, for example, for time-of-flight imaging or for rendering of light echoes. This work was published in a regular journal paper at ACM Transactions on Graphics [2] and was presented at ACM SIGGRAPH 2014.

Figure 6 shows how a laser beam is continuously refracted in a sugar solution with varying concentration. The setup consists of a glass tank with fresh water intermixed with one droplet of a 2% eosin solution to increase scattering of light. Afterward, 2 kg of sugar cubes were dissolved on the bottom of the tank. By diffusion and gravity, a concentration profile formed naturally over a time span of 24 h with a decreasing sugar concentration from the bottom toward the top. The index of refraction varies between $n = 1.42$ (ca. 50% concentration) on the bottom and $n = 1.33$ (fresh water) in a height of about 4 cm and above. Figure 6(a) shows a photograph of the real-world experiment. A green 1 mW laser pointer emits a collimated beam of light that enters the tank horizontally and gets continuously refracted in the sugar solution. On the bottom of the tank, the beam gets partially reflected and is again bent downward. The virtual experiment, in Figure 6(b), closely reproduces the trajectory of the beam, its reflections on the back side of the tank, and the scattering of light in the fluid.

Figure 7 shows renderings of a cup with a spatially varying index of refraction due to the mixing of melted glass and jadeite. In Figure 7(a), light transport is solved according to the standard radiative transfer equation, which leads to incorrect conservation of energy. In Figure 7(b), light transport is described by the extended refractive radiative transfer equation, which accounts for correct conservation of energy.

9 Sort First Parallel Volume Rendering

The paper *Sort First Parallel Volume Rendering* [12] presents a distributed rendering algorithm that is based on an image-space decomposition with dynamic load balancing for GPU-based cluster systems. The paper presents a bricking approach that offers improved data scalability, which is a weak point of sort first decompositions, especially for GPU-based rendering. Furthermore, a novel caching strategy is described that exploits frame-to-frame coherence to preload data that is close to the current frustum of each render node. Depending on the viewpoint, the computational load can strongly vary between the render nodes. Therefore, a dynamic load balancing algorithm is introduced that reorganizes the image decomposition based on a per-pixel cost estimation. It is also shown that sort first partitioning is prefer-

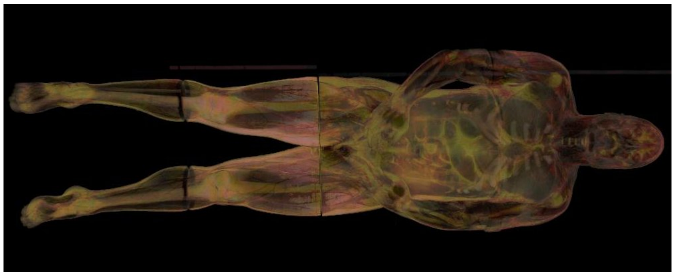

Fig. 8: Parallel volume rendering of the Visible Male data set ($2048 \times 1024 \times 1878$) with a low opacity transfer function.

able to parallelize algorithms that depend on ray coherence such as visibility culling or volumetric shadows. The paper was published in a regular journal paper at IEEE Transactions on Visualization and Computer Graphics [12] in 2011.

Figure 8 shows a parallel volume rendering of the Visible Male data set at interactive frame rates on a distributed GPU-cluster (32 nodes) with dynamic load balancing in image space.

10 3D Reconstruction of Astronomical Nebulae

The papers *Visualization of Astronomical Nebulae via Distributed Multi-GPU Compressed Sensing Tomography* [13] and *Interactive Visualization and Simulation of Astronomical Nebulae* [14] present algorithms to reconstruct, simulate, and visualize 3D models of astronomical objects like supernovae or emission nebulae. Most important, a compressed sensing algorithm is introduced to automatically reconstruct a volumetric model of astronomical nebulae from a single telescope image. The method exploits the approximate symmetries of these objects to estimate the missing information from other directions. To avoid overly symmetric results, the technique further introduces artificial jittering and enforces that the reconstructed model exactly matches the view from the input image. This combination of building blocks provides plausible models that can be visualized interactively with direct volume rendering. The optimization problem is solved numerically with a distributed approach on a GPU cluster. This technique was published in a conference paper at IEEE SciVis 2012 [13] and the 3D models of the nebulae were integrated in the DigiStar framework for planetarium shows by Evans & Sutherland.

The Butterfly Nebula, or M2-9, is an example of a bipolar planetary nebula whose structure is roughly described by an approximate axisymmetry as shown in Figure 9(a). Even though the assumed symmetry is only approximate, most details are clearly visible in the reconstructed volume in Figure 9(b).

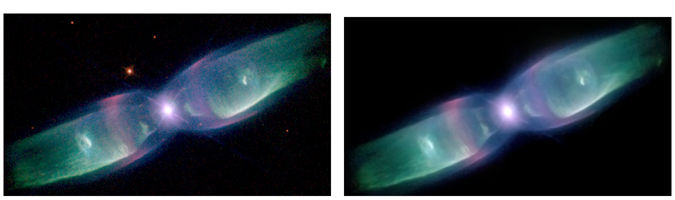

(a) (b)

Fig. 9: Planetary nebula M2-9. (a) Input image from Hubble Space Telescope. (b) High-resolution 3D visualization that closely resembles the original image when rendered from the same viewpoint. *Original image: Bruce Balick (University of Washington), Vincent Icke (Leiden University, The Netherlands), Garrelt Mellema (Stockholm University), and NASA.*

References

[1] M. Ament. *Computational Visualization of Scalar Fields*. PhD thesis, University of Stuttgart, 2014.

[2] M. Ament, C. Bergmann, and D. Weiskopf. Refractive radiative transfer equation. *ACM Transactions on Graphics*, 33(2):17:1–17:22, 2014.

[3] M. Ament and C. Dachsbacher. Anisotropic ambient volume shading. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):1015–1024, 2016.

[4] M. Ament, S. Frey, C. Müller, S. Grottel, T. Ertl, and D. Weiskopf. GPU-accelerated visualization. In E. W. Bethel, H. Childs, and C. Hansen, editors, *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*, pages 223–260. Chapman and Hall/CRC, 2012.

[5] M. Ament, F. Sadlo, C. Dachsbacher, and D. Weiskopf. Low-pass filtered volumetric shadows. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2437–2446, 2014.

[6] M. Ament, F. Sadlo, C. Dachsbacher, and D. Weiskopf. 2015 cover image (2nd place): Vortex cascade. *Computer Graphics Forum*, 2015.

[7] M. Ament, F. Sadlo, and D. Weiskopf. Ambient volume scattering. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2936–2945, 2013.

[8] M. Ament, F. Sadlo, and D. Weiskopf. 2014 cover image: Supernova. *Computer Graphics Forum*, 33(1):305–306, 2014.

[9] M. Ament, D. Weiskopf, and H. Carr. Direct interval volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1505–1514, 2010.

[10] M. Ament, D. Weiskopf, and C. Dachsbacher. Ambient volume illumination. *IEEE Computing in Science and Engineering*, 18(12), 2016.

[11] M. Ament, T. Zirr, and C. Dachsbacher. Extinction-optimized volume illumination. *IEEE Transactions on Visualization and Computer Graphics*, to appear, 2016.

[12] B. Moloney, M. Ament, D. Weiskopf, and T. Möller. Sort first parallel volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 17(8):1164–1177, 2011.

[13] S. Wenger, M. Ament, S. Guthe, D. Lorenz, A. Tillmann, D. Weiskopf, and M. Magnor. Visualization of astronomical nebulae via distributed multi-GPU compressed sensing tomography. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2188–2197, 2012.

[14] S. Wenger, M. Ament, W. Steffen, N. Koning, D. Weiskopf, and M. Magnor. Interactive visualization and simulation of astronomical nebulae. *IEEE Computing in Science and Engineering*, 14(3):78–87, 2012.