Report of the
Dagstuhl Seminar on Rendering

June, 10–14, 1996

Organized by

Patrick M. Hanrahan, Stanford University

Heinrich Müller, Universität Dortmund

Claude Puech, Université Joseph Fourier, Grenoble

Rendering is a special field of computer graphics. Its subject is the synthesis of images or image sequences from the description of a three-dimensional scene. The range of this topic goes from on-line image synthesis, supported by special purpose hardware, to detailed simulations of optical models of light. Examples are the realistic display of a car body designed in a CAD system using graphics workstations equipped with shading hardware, or the simulation of a room which considers the correct technical parameters of the light sources installed. Other applications are in visual simulation systems for the training of pilots or drivers of cars or ships for which optical effects in participating media are of particular importance.

Face to face with these economically relevant applications basic research problems emerge from the demand of fast image synthesis in applications on one side, and the inherent computational complexity of this problem on the other side. A typical task is the development of approximate solutions allowing the calculation of correct optical effects with sufficient accuracy, but with practical efficiency. Another emphasis lies on the development of solutions in massively parallel environments. In the last ten years, significant progress has been made, but still there remains considerable work to be done. The list of participants of this Dagstuhl Seminar comprehends many researchers significant in this field.

One idea of this Dagstuhl Seminar was to exchange the requirements of systems, interfaces, and efficient algorithmic solutions to be developed. Another goal of the seminar was to provide an opportunity for discussing ideas and work in progress. International conferences with their densely packed schedules leave practically no room for this sort of scientific exchange. The reactions of the participants after the seminar showed that these aims were evidently achieved. There were unexpected results and new ideas that caused extraordinarily intensive and emotional discussions, and the excellent session about Monte-Carlo techniques was a unique event.

The particular atmosphere of Dagstuhl was additionally stimulative, and obliging staff, good food, fine weather, and the heavily used new sporting facilities of Dagstuhl certainly contributed to the harmonic atmosphere among the about fourty participants from twelve countries. There is a requirement of an event of the type of Dagstuhl Seminars, and it was the broad opinion that a new seminar on rendering should take place in the future.

Participants

Michael Bender, Universität Kaiserslautern, Germany
Christian Bohn, GMD-IMK, Germany
Alan Chalmers, University of Bristol, Great Britain
Michael Cohen, Microsoft Research, USA
Philip Dutré, Katholieke Universiteit Leuven, Belgium
Dieter Fellner, Universität Bonn, Germany
Alain Fournier, University of British Columbia, Canada
Bernd Fröhlich, Stanford University, USA
Robert Garmann, Universität Dortmund, Germany
Markus Groß, ETH Zentrum, Switzerland
Hans Hagen, Universität Kaiserslautern, Germany
Philip Jacob, Universität Kaiserslautern, Germany
Frederik W. Jansen, Delft University of Technology, The Netherlands
Henrik Wann Jensen, Technical University of Denmark, Denmark
Alexander Keller, Universität Kaiserslautern, Germany
Eric P. Lafortune, Katholieke Universiteit Leuven, Belgium
Dani Lischinski, University of Washington, USA
Nelson Max, University of California, Lawrence Livermore National Lab., USA
Heinrich Müller, Universität Dortmund, Germany
Stefan Müller, Fraunhofer Institut für Graphische Datenverarbeitung, Germany
Laszlo Neumann, Budapest, Hungary
Gregory M. Nielson, Arizona State University, USA
Georg Pietrek, Universität Dortmund, Germany
Claude Puech, Universite Joseph Fourier, France
Xavier Pueyo, Universita de Girona, Spain
Werner Purgathofer, TU Wien, Austria
Georgios Sakas, Fraunhofer Institut für Graphische Datenverarbeitung, Germany
David Salesin, University of Washington, USA
Stephan Schäfer, Universität Bonn, Germany
Peter Schröder, California Institute of Technology, USA
Mark Segal, Silicon Graphics, USA
Hans-Peter Seidel, Universität Erlangen, Germany
Philipp Slusallek, Universität Erlangen, Germany
Marc Stamminger, Universität Erlangen, Germany
Wolfgang Stürzlinger, Johannes Kepler University, Austria
Laszlo Szirmay-Kalosm, Technical University of Budapest, Hungary
Daniela Tost, Universitat Politecnica de Catalunya, Spain

Table of Contents

Peter Schröder
Building your own wavelets at home

Bernd Fröhlich, Georg Pietrek, Richard Bartels
Hierarchical Continuous Basis Functions for Radiosity

Alexander Keller
Quasi-Monte Carlo Methods for Photorealistic Image Synthesis

Dieter Fellner
An Engineering Approach to Graphics Architectures

Eric Lafortune
Monte Carlo Techniques and Optimisations for Physically Based Rendering

Nelson Max
Hierarchical Rendering of Trees from Precomputed Multi-Layer Z Buffer

Xavier Pueyo, Gonzalo Besuiewsky, Francesc Castro
Making Global Monte Carlo Algorithms Usable

Dani Lischinski
Hierarchical Image Caching for Accelerated Walkthroughs of Complex Environments

Dani Tost, Anna Puig, Isabel Navazo
Features enhancement in volume rendering – A case study: celebral blood vessels visualization

László Neumann, Attila Neumann
A New BRDF Model with Fast Importance Sampling

Alain Fournier, Bob Lewis
Wavelet Radiative Transport and Surface Interaction

Henrik Wann Jensen
Global Illumination using Photon Maps

Peter Neugebauer, Georgios Sakas
"Virtual Cloning" – Creating Accurate Geometric Models of 3D Objects Using Range Images

Coherence Improving Techniques for Parallel Rendering

Frederik W. Jansen
Delft University of Technology

For a parallel implementation of ray tracing and ray-tracing-based radiosity we exploit a demand-driven scheduling strategy for the primary and shadow rays because these show significant object coherence, and a data-parallel scheduling approach for the secondary, and less coherent rays. By taking the data parallel tasks as a basis load and the demand-driven tasks as an additional load, a very efficient and scalable algorithm can be achieved.

However, it turns out that for large and complex scenes the data parallel processing is dominant, resulting in an unbalanced load near the end of the processing (when all the demand driven tasks are issued) and in a long latency for the last rays to finish.

Currently, we try to improve the performance of the system by avoiding excessive sampling of secondary rays by replacing the explicit tracing with an approximated value derived from a radiosity preprocessing or an illuminance caching scheme. In addition we explore the use of environment maps and grouping to replace the explicit sampling of complex environments by a look-up in one of these image-based representations. Combining these techniques with illuminance caching, hopefully gives us the desired scalable algorithm to render very large and complex scenes.

Ray Tracing of Parametric Surfaces: Bézier Clipping, Chebyshev Boxing, and Bounding Volume Hierarchy – A Critical Comparison with New Results

S. Campagna, Ph. Slusallek, H.-P. Seidel
Universität Erlangen, Germany

Parametric surfaces are a powerful and popular modeling tool in computer graphics and computer aided design. Ray tracing is a versatile and popular rendering technique. There is therefore a strong incentive in developing fast, accurate and reliable algorithms to solve the ray patch intersection problem.

In this talk we discuss and compare three recent geometric algorithms for solving the ray patch intersection problem, namely Bézier Clipping (Nishita et al. '90), Chebyshev Boxing (Fournier, Buchanan '94), and Bounding Volume Hierarchy (Campagna'95). We focus on issues such as speed, robustness, and memory requirements and discuss how various user selected parameters affect the performance of the algorithms. We find that Bézier Clipping may be preferred as the initial default choice for producing high quality pictures without artifacts in reasonable time and with little memory requirements, while the two other methods may be advantageous if the same object has to be rendered multiple times.

All three algorithms have been implemented within the 'Vision' rendering system, and various examples and timings illustrate our results.

[FB94] A. Fournier and J. Buchanan. Chebyshev polynomials for boxing and intersections of parametric curves and surfaces. *Computer Graphics Forum* 13(3), (Proc. EUROGRAPHICS '94), 127 – 142, 1994.

[**NSK90**] T. Nishita, T. Sederberg, and M. Kakimoto. Ray tracing trimmed rational surface patches. *Computer Graphics*, 24(4), (Proc. SIGGRAPH '90), 337 – 345, August 1990.

[**Cam95**] S. Campagna. Vergleich und Erweiterung von Verfahren für die Schnittpunktberechnung von Strahlen mit Polynomflächen. Diplomarbeit. Universität Erlangen, Computer Graphics Group, July 1995 (in German).

The Lumigraph

Michael F. Cohen, Steven Gortler, Radek Grzeszczuk, Riek Szeleski
Microsoft Research, Richmond WA, USA

The lumigraph is a 4D function representing the complete flow of light out from a bounded object (or conversely into an empty region). Given a lumigraph, one can render an image of an object quickly from any point outside the object's convex hull. This talk describes the lumigraph function, a choice of parameterization based on ray intersections with two planes, discretization and approximation as a linear sum of basis functions, and finally reconstruction of images taking advantage of texture mapping hardware. The talk also describes adaptation of the 4D basis function shapes based on an approximate geometric model of the object.

A system for capturing the light leaving a real world object is shown utilizing a hand held video camera. The captured image data is "rebinned" into the lumigraph structure with a linear time scattered data interpolation method. Images are finally created from arbitrary virtual cameras in real-time.

Extensions of the current implementation were discussed for use in outdoor scenes or scenes with multiple objects.

On the Complexity of Ray-Shooting

László Szirmay-Kalos, Márton Gábor
Technical University of Budapest

This paper examines ray-shooting from both worst-case and average case complexity point of views. It introduces an algorithm, called the complementer plane algorithm, that can solve ray-shooting for any convex objects in logarithmic time in the worst case. The storage complexity of this method is $O(n^8)$. It is proven that logarithmic time is optimal for the ray-shooting problems. Based on the observation that logarithmic search requires an ordered set, it is shown that logarithmic ray-shooting algorithms need at least $\Omega(n^4)$ storage space, which makes this approach not very suitable for practical implementations.

Practical implementations tend to use heuristic methods that aim at average-case optimization. Thus worst-case complementer analysis cannot evaluate them properly. We have to apply average-case complexity measures which determine the expected running time and storage. To do so a probabilistic model of the object configuration space is required. This paper proposes a model which assumes that the space contains spheres of the same size and sphere centers are uniformly distributed. It is demonstrated that heuristic ray-tracing acceleration techniques will run in constant time under this model, including regular grid, ray coherence and Voronoi-diagram based space partitioning.

Hardware-Accelerated Rendering of Curved Surfaces

Mark Segal
Silicon Graphics Computer Systems

Current graphics hardware architectures are designed to draw polygons quickly and effectively. Curved surfaces are typically drawn by tesselating them on the host processor and sending the resulting triangles to the graphics subsystem. While this technique is effective, it has drawbacks. Dynamic tesselation is relatively slow, often leaving the graphics pipeline on which it is being carried out idle. Many triangles must be used for a given surface to achieve good rendering quality, but too many triangles can be slow to transfer to and process on the graphics hardware.

We describe the functional components of the hardware rendering pipeline in use today with an eye towards how it could be employed to render curved surfaces directly or could be modified to do so. We examine three classes of surface representations: implicit, parametric, and limit, and discuss the opportunities and challenges for their accelerated rendering in graphics hardware.

Decoupling Transport and Reflection for Radiance Computations

Marc Stamminger, Philipp Slusallek, Hans-Peter Seidel
University of Erlangen

Although significant success has been made in the improvement of radiosity algorithms, the methods to solve the general global illumination problem with arbitrarily reflecting surfaces still require unreasonable computing times. In this talk, we try to identify the reasons and present the concept of computing transport and reflection operator separately to avoid this behavior.

We describe an extension to the well-known hierarchical radiosity algorithm, which can also handle non-diffuse surfaces and uses distinct representations of the transport and reflection operator. Several possibilities how the operator can be computed and their impact on the solution are discussed. Finally, we present first results of our implementation.

Realtime Rendering and Radiosity – Applications and Problems

Stefan Müller
Fraunhofer - IGD, Darmstadt

In Virtual Reality and its application, we have still many problems which are related to the field of rendering. In this talk, some applications are described and the rendering problems are emphasized. In the field of virtual prototyping there is an important need of hardware rendering of non-polygonal objects (like spheres, cylinders etc.). Automatic LOD-generation and scene-reduction is also needed. For architectural scenes, radiosity can often not be used as a global illumination algorithm; depending on the perception of the user, we use radiosity locally and decide meshing and adaptation algorithms interactively.

An important issue for the future will be the combination of real and synthetic images. A few examples for reconstruction of objects out of video images and augmented reality are also shown.

Parallel Processing for Photo-Realistic Archaeological Site Visualization

Alan Chalmers
University of Bristol, UK

Recent developments in computer graphics enable virtual environments to be created and viewed photo-realistically on a computer. Archaeologists want to use this technology to visualize sites and develop hypotheses concerning how these sites may have been utilized by our ancestors. One example of the need for such a system is in order to appreciate the meaning of light in medieval environments. Medieval pottery was often vibrantly decorated: Did medieval people simply enjoy bright colours, or was such colouring essential to distinguish different pots in the poorly lit medieval homes?

This talk discusses the multi-disciplinary approach that we feel must be taken in order to achieve the desired level of realism archaeologists will need in order to assist them in answering these questions. The particle tracing method is selected to compute the illumination within the environments. This method is particularly suitable as it is able to cope with participating media, a fundamental component of any historical environment, in a straightforward manner. To accurately simulate such environmental factors as smoke, fog, dust, and flame, computational fluid dynamics and computer vision techniques have been incorporated. The talk explains the important role computer vision will have for future of computer graphics.

The computational requirements for achieving the desired level of realism is significant. Parallel processing offers the potential of producing images in reasonable time. The talk concludes with details of the SAMD model of parallelism which has been developed to efficiently implement our archaeological system, INSITE, on a multiprocessor machine.

Optimizing the Local Pass

Wolfgang Stürzlinger
University of Linz, Austria

Recently two pass approaches have been introduced in photorealistic image synthesis. Most global illumination methods cannot produce an image without visual artifacts due to the unknown viewing parameters and interpolation errors. Therefore a second so-called local pass is used to compute the illumination of each pixel of the image. This is done by computing the local illumination for the corresponding point on the visible surface using the global illumination solution.

An exact method is presented based on the projection of polygons on the hemisphere. Computing the results is done in three dimensions, but the used data structure is two-dimensional, thereby avoiding the need to process (curved) edges on the hemisphere.

A faster alternative to calculate the local illumination was also presented. It uses importance sampling to re-evaluate visibility along the links provided by the hierarchical radiosity solution.

This method is able to achieve a good approximation of the exact value with a low number of visibility rays.

Multiresolution Modelling and Image Generation

Markus Groß
Swiss Federal Institute of Technology (ETH Zürich)

Multiresolution Analysis and Wavelets provide many useful properties, which enable to build fancy rendering, approximation and compression methods. Especially the finite support of the basis allows to define various weighting functions which control the accuracy of the wavelet expansion in terms of spatially variant levels-of-detail. In this talk, two different applications of wavelets in computer graphics were presented. The first one comprised an adaptive mesh refinement and compression strategy, where the overall approximation error is governed by local spectral estimates derived from a WT. The method builds up a quadtree representation by progressively removing vertices from the mesh. The corresponding criteria are determined by analyzing individual difference spaces and by considering topological constraints.

The second example encompassed the fast computation of line integrals through volume data sets. Starting with an initial 3D WT of the data, the method computes a set of footpoints of the basis functions which can be splatted into the frame buffer. Progressivity, self-similarity and compression gain are features of this algorithm, which is suitable for networked applications.

Towards an Open Rendering Kernel for Image Synthesis

Philipp Slusallek
IMMD IX - Universität Erlangen

In order to use realistic image synthesis successfully in research and development as well as in commercial products, two important prerequisites have to be fulfilled. First of all, good, accurate, robust, and fast algorithms are required. Impressive progress has been made in this respect during the last years. The second step is the creation of a suitable and general software architecture that offers an environment into which these rendering algorithms can be integrated.

In this presentation, we develop an architecture that consists of a small but flexible rendering kernel. This kernel provides a general framework for rendering algorithms and defines suitable interfaces for specific aspects of rendering, like reflection or emission. Algorithms for a particular aspect of the rendering process can then be plugged into the kernel in order to implement a special rendering strategy. The benefits of this approach are demonstrated with several applications.

Hierarchical Radiosity - An Analysis of Computational Complexity for a Simple Scene

Robert Garmann
Universität Dortmund

The hierarchical radiosity algorithm is an efficient algorithm for the simulation of light. Hanrahan et al. describe the initialization and the refinement of links between the scene's patches based upon a user-specified error parameter ε. They state that the number of links is linear in the number of patches if ε is assumed to be a constant.

We show a more natural result for a very simple scene consisting of two parallel one-dimensional equally sized patches. The proof uses the formfactor as an oracle of the error involved with one link. Based on the assumption that the geometry is constant and ε approaches 0 in a multigridding procedure we show that the algorithm generates $N = O(\varepsilon^{-1})$ subpatches and $L = \Theta(N^2)$ links.

A Mathematical Framework for Global Illumination

Philip Dutré
Dept. of Computer Science, Katholieke Universiteit Leuven, Belgium

The last few years, some papers have been describing the use of dual equations in global illumination algorithms. The potential equation can be derived as the dual to the rendering equation. Computing the potential equation with a Monte Carlo method, results in particle tracing which is the "dual algorithm" to ray tracing.

Both dual equations can be combined in a single function which is called the global reflection distribution function (GRDF) (Lafortune '94).

Series expansion of the dual equations results in a "triangle" of light – potential interactions, which can be used to derive new algorithms, or to discover already proposed ones.

Efficiently Representing the Radiosity Kernel through Learning

Christian A. Bohn
Dept. Visualization and Media Systems Design
German National Research Center for Information of Technology
Schloß Birlinghoven, Sankt Augustin, FR Germany

A new approach for approximating the radiosity kernel by a discrete set of basis functions is presented. The algorithm is characterized by selecting samples from the geometry definition and iteratively creates a functional model instantiated by a set of Gaussian basis functions. These are supported over the whole environment and thus, surfaces are not considered separately. Together with the implicit clustering algorithm provided by the applied learning scheme, the algorithm accounts ideally for coherence in the global kernel function.

On one hand, this leads to a very sparse representation of the kernel. On the other hand, by avoiding the creation of initial basis functions for separate pairs of surfaces, the method is

capable of calculating even huge geometries to a desired accuracy with a proportional amount of computing resources.

Recent results from the field of artificial neural networks (the Growing Cell Structures) are extended for the presented learning algorithm. This work is done in Flatland, but there are no methodical constraints which bound the application to two dimensions.

Building Your Own Wavelets At Home

Peter Schröder, Caltech USA
Wim Sweldens, Lucent Technologies, USA

The lifting scheme is a construction technique for wavelets which is performed entirely in the spatial domain. This allows us to generalize many traditional constructions to the second generation setting: irregular sample location, weighted measures, intervals, etc. The construction is based on a simple factorization of the filter sequence which can be understood as a sequence of Prediction and Update steps. The former compute new values in between old values, for example, as in subdivision as known in CAGD. The latter ensures the preservation of moments, which is important for stability reasons. The method is simple to implement, can reecute in place, is trivial to invert by simply running the code "backwards" and exposes maximal parallelism. However, it is purely algorithmic and analytic properties need to be checked in each case.

Hierarchical Continuous Basis Functions for Radiosity

Bernd Fröhlich, GMD, St. Augustin
Georg Pietrek, Universität Dortmund
Richard Bartels, C G L, University of Waterloo

We showed how wavelets with overlapping support can be used for the hierarchical radiosity approach. We investigated two wavelet families: Daubechies wavelets adapted to the interval and B-spline wavelets. Daubechies' construction has the advantage to be orthenormal whereas B-splines have well-known approximation properties. We extended both families of wavelets to represent functions at coarser scales than the standard construction. This is important for radiosity since most energy transfers can be computed accurately enough with coarse representations of the radiosity function. The mathematical framework for the use of this type of basis functions is developed and first results were shown.

Quasi-Monte Carlo Methods for Photorealistic Image Synthesis

Alexander Keller
University of Kaiserslautern, Germany

The global illumination problem in computer graphics is described by a second kind Fredholm integral equation. We present the Quasi-Random Walk algorithm for calculating functionals of the solution of the integral equation. This consistent, deterministic algorithm uses low discrepancy sequences and outperforms stochastic random walk techniques. We also showed the efficiency of low discrepancy quadrature applied to other problems of image synthesis such as the local pass and form factors. Finally we provide a simple construction rule for such algorithms.

An Engineering Approach to Graphics Architectures

Dieter Fellner
Universität Bonn

We present an object-oriented software architecture for a 3D rendering environment which drastically improves the programmer's productivity, and, most importantly, consists of building blocks that lend themselves to customization thus making 3D image synthesis more accessible.

The rendering platform MRT (for Minimal Rendering Tool) has been designed to

- bridge the gap between modeling and rendering,
- provide suitable rendering quality (from consistent handling of approximate shading to ray tracing, radiosity, and to hybrid methods),
- provide platform independence,
- increase productivity, and to
- make rendering technology generally available.

Experiences with our (inhomogeneous) user population from undergraduate students to users in industry prove that the system meets its design goal of being highly customizable and extendable. Furthermore, it serves as a compact testbed for various rendering aspects as well as for new algorithms 'out side' of the classical rendering domain. This is supported by an industry cooperation for the development of a prototype to simulate the distribution of radio waves in 3D urban environments which could be completed by one of our students within three weeks. The incredible short development time (considering that we started from scratch) in combination with the fact that the prototype was significantly faster than what has been available before was the beginning of a very productive research cooperation.

Further information (including details on how to get the software by anonymous ftp) can be found at `http://hyperg.cs.uni-bonn.de/CompGraph`.

Monte Carlo Techniques and Optimisations for Physically Based Rendering

Eric Lafortune
Dept. of Computer Science, K.U. Leuven, Belgium
Program of Computer Graphics, Cornell University, USA

In spite of constant progress on sophisticated deterministic solution methods for the global illumination problem, Monte Carlo techniques still play an important role, e.g. for final gather passes. I present an overview of classical Monte Carlo techniques and their application to physically based rendering. I recall the basic principles and then discuss the influence of taking more samples, stratifying them, using antithetic variates, control variates, importance sampling, adaptive sampling and Russian roulette. Graphs illustrate how variance reduction techniques are nothing else than transformations of the original integral. The resulting functions are hopefully flatter, so that they are easier to integrate. Specific results obtained from tests with the optimisations, for a set of test scenes rendered with path tracing and bidirectional path tracing, indicate the importance of stratified sampling and of importance sampling for these algorithms.

Hierarchical Rendering of Trees from Precomputed Multi-Layer Z Buffer

Nelson Max
University of California, Davis/Livermore

Recent work in image-based rendering takes advantage of surface coherence and image morphing for reprojecting a precomputed or photographed image onto a new view. However, trees are complex volume objects where every pixel may be on a different leaf at a different depth so morphing is not appropriate. Here I reproject each pixel from a z buffer separately, using the depth value in the reprojection. Images are precomputed for twigs and branches at various levels in the hierarchical structure of a tree, and adaptively combined, depending on the position of the new viewpoint. I used 22 orthogonal precomputed views of each object, on the unit sphere of viewing directions. The precomputed images contain multiple z levels to avoid missing pixels in the reconstruction, subpixel masks for anti-aliasing, and colors and normals for shading and shadowing after reprojection. I also proposed, but have not yet implemented, a multi-layer, multi-resolution shadow buffer, for more accurate shadowing, including penumbras, which contain more detail near the viewpoint, again adaptively constructed from the image hierarchy.

Making Global Monte Carlo Algorithms Usable

Xavier Pueyo, Gonzalo Besuiewsky, Francesc Castro
Universitat de Girona

Global Monte Carlo (GMC) algorithms have been used by several authors and have proven their interest for certain types of complex scenes. In the talk we have presented an overview of GMC algorithms as well as of previous improvements proposed to the basic approach. After this introduction, we have presented two alternatives of adaptive solution.

The first approach comes from the evidency that if we use a unique density of lines for all the scene, either it will be too precise (so, unnecessarily expensive) for certain subscenes, or will be under the required precision for some other subscenes. To solve this, we have proposed the use of a hierarchy of densities of lines. We define a hierarchy of clusters (spheres or parallelepipeds) and a density of lines for each cluster. For modeling the interaction between two surfaces into a given cluster, we may use the lines of this cluster as well as the lines of uper levels of the hierarchy.

Another solution has been presented based on a first division of the space occupied by the scene into tubes which are afterwards adaptively subdivided. If we assume a spherical cluster containing the complete scene, its surface may be discretised into polygons of equal area. Each pair will then define a tube. Next operation is to estimate the "complexity" of each tube in order to determine how many samples must be taken into it and to decide to proceed, or not, to subdivide the tube. Special attention must be paid for maintaining the uniformity of the density of lines.

Hierarchical Image Caching for Accelerated Walkthroughs of Complex Environments

Dani Lischinski
University of Washington

We present a new method that utilizes path coherence to accelerate walkthroughs of geometrically complex static scenes. As a preprocessing step, our method constructs a BSP-tree that hierarchically partitions the geometric primitives in the scene. In the course of a walkthrough, images of nodes at various levels of the hierarchy are cached for reuse in subsequent frames. A cached image is reused by texture-mapping it onto a single quadrilateral that is drawn instead of the geometry contained in the corresponding node. Visual artifacts are kept under control by using an error metric that measures the discrepancy between the appearance of the geometry contained in a node and the cached image. The new method is shown to achieve speedups of an order of magnitude for walkthroughs of a complex outdoor scene, with little or no loss in rendering quality.

Features Enhancement in Volume Rendering – A Case Study: Celebral Blood Vessels Visualization

Dani Tost, Anna Puig, Isabel Navazo
Polytechnical University Calabria, Barcelona

The reconstruction of the celebral vascular structure and the use of accurate visualization techniques can greatly improve the conditions of the diagnosis of vascular pathologies such as hemonagies and emblia caused by aneurysms and stenoses. Currently, the registration of the brain and vascular structures is done using MR and MRA images which present a lot of noise and artifact and fail at enhancing small vessels, stationary flux and real diameters. A segmentation step based on photometric properties as well as the knowledge of the tubular structure of the vessels is needed. After the segmentation, two strategies can be adopted for the visualization: (i) extracting the inner surface of the vessels and (ii) projecting directly the volumetric model. The first strategy presents two major drawbacks. First, current surface extraction methods like marching cubes are not suitable for vessels which exhibit a very complicate and narrow shape. A marching cube extraction of a vessel surface gives a surface model composed of a large amount of tiny polygons and does not guarantee continuity. The second drawback ist that surface visualization does not allow the integration of the vessels with the brain.

Direct visualization, either by using ray casting or a BTF or FTB projection, has proved to be the more suitable visualization technique. The shading model to be used considers that the structures are homogeneous in the interior (emission and absorption) and have surface scattering in the boundary.

Different issues on the improvement of direct visualization are analyzed including: (i) the use of compressed data structures to speed up the direct access to the relevant samples, (ii) the preprocessing of view independent parameters such as surface normals, (iii) the reduction of the projection using a coherent splatting method.

Future work will include the construction of a symbolic model from which a simple soft surface can be reconstructed and used in a mixed visualization.

A New BRDF Model with Fast Importance Sampling

László Neumann, Attila Neumann
Budapest, Ungarn

We introduce a physically plausible mathematical model for a large class of BRDFs. The new model can be defined in a user-friendly way and it is as simple as the well-known Phong model, without any of its disadvantages.

It gives a good visual approximation of a lot of practical materials: metals, plastics, ceramics, retro-reflective paints, anisotropic materials, etc. The new model includes the ideal diffuse model and ideal mirror as well, reflecting all of the incoming energy.

Because of its illustrative properties it can be used easily in most commercial software and because of its low computational cost it is practical for VR.

The model is based on a special basic BRDF definition, which fulfills the requirements of reciprocity and of energy conservation. We then construct the new model from this basic BRDFs with different weight functions. Finally we can – using 1D and 2D uniform random points – quickly generate directions following the distribution of reflected energy.

Wavelet Radiative Transport and Surface Interaction

Alain Fournier, Bob Lewis
University of British Columbia

We describe how light flux is represented and operated on in the context of a "light-driven" global illumination approach. The basic approach consists in carrying illumination through virtual walls in the environment. These walls are determined adaptively and recursively by an octree data structure.

The light flux on these walls are described by 4D wavelet bases (two dimensions in space and two dimensions in directions). After an introduction on the basic characteristics of the global illumination computations, which we call "Lucifer", we concentrate on the operations necessary to carry out the propagation of the light flux through the octree cells, express the blocking by objects in the environment and compute the reflected/transmitted light.

Propagation is computed through transfer coefficients between the wavelet basis functions of the source wall and the destination wall. Computing these coefficients take a 4D integration, and we compare various methods (Monte Carlo, Romberg, trapezoid, quasi-Monte Carlo) as for the accuracy/speed trade-off. Blocking is computed through a simple adaptive algorithm which determines the level of details of the light flux on the destination walls. Finally reflection/transmission is computed in a way similar to the propagation, where the BRDF is expressed as the coupling coefficients between the wavelet basis in the incident flux and the wavelet basis of the reflected flux.

Global Illumination using Photon Maps

Henrik Wann Jensen

Current state of the art techniques in global illumination are two-pass methods in which the first pass is a radiosity like algorithm that creates an approximate global illumination solution. In the second pass this approximation is visualized using an optimized Monte Carlo ray tracer. This scheme works very well in most scenes but as models become increasingly complex having millions of primitives, procedural objects and glossy reflection the cost of using radiosity becomes prohibitive. This is mainly due to the fact that storing illumination within a tessellated representation of the geometry uses too much memory.

In this talk a two-pass global illumination method is presented which uses the photon map to represent illumination within a model. The photon map is created by emitting a large number of photons (packets of energy) from the light sources into the scene. Each photon is traced through the scene using photon tracing (similar to path tracing). Every time a photon hits a non-specular surface it is stored in the photon map. The result is a large number of photons stored within the model approximating the incoming flux at the surfaces. The photon map can be used to provide radiance estimates at any given surface position, to generate optimized sampling direction in a Monte Carlo ray tracer, to reduce the number of shadow rays (using shadow photons) and to provide improved control variates.

The method presented uses two photon maps: One high resolution photon map representing caustics that are visualized directly and one lower resolution global photon map which is used to reduce the number of reflections traced and to generate optimized sampling directions in the Monte Carlo ray tracer.

Results are presented which demonstrate global illumination in scenes containing procedural objects and surfaces with diffuse and glossy reflection models. The implementation is also compared with the Radiance program.

“Virtual Cloning“
Creating Accurate Geometric Models of 3D Objects
Using Range Images

Peter Neugebauer, Georgios Sakas
Fraunhofer - IGD, Darmstadt

The major efforts in Computer Graphics during the recent decades has been directed towards the generation of (realistic) images. This process of converting semantically high object descriptions (i.e. geometric object moodels) to pixels, i.e. semantically lower ”signals”, is inherently associated with a loss of information.

However, in order to generate realistic renderings of real-life objects such as buildings, statues, trees, cars, tea-pots etc. a realistic model is required first. This model has either to be generated (i.e. manual modelling, procedural generation etc.) or to be captured from the real-life original. Texture mapping is a good understood example of "realistic input for realistic output" working paradigm.

Our research target goes beyond this classic application area. We are focusing on capturing ALL aspects of a real-life object in order to generate a high-quality computer graphics model suitable for the usual applications. Therefore we capture:

1. geometry
2. texture
3. surface reflectance properties

For capturing geometry we use a structure light range image camera. However, our approach works identically with any other active or passive range image scanner (laser scanner, stereo, Moiree etc.). The desired object is scanned from several directions. In the second step all of the range images are simultaneously registered, so that their relative orientation in 3D space is defined. Our approach employs optimization best-matching of the overlapping regions. Due to the simultaneous registration of ALL available images, error accumulation is avoided and a sub-pixel precision is achieved.

The second and most important novelty of our approach is that before we generate polygonal descriptions we create first a volumetric model (topogram) of the reconstructed object. We call our technique "sculpturing" since it works very similar to the artistic sculpturing method: Starting with a unit solid block of free-defined resolution (typically 256^3), each range image can only extract "material" form the volume. Thus ambiguities, self-shadowing etc. are automatically solved and all types of models, convex and concave, can be generated.

In the last step we use marching cubes to create the polygonal representation. By low-pass filtering the original volume prior to marching cubes, we have a very elegant, easy-to-implement and mathematically solid method for creating level-of-detail models.

Our approach has been demonstrated in modeling a human-size statue (1.82 m) in a Bavarian Museum. For reconstructing the extremely complicated statue having holes, concavities etc. as many as 180 images from all directions and of different scales (general views as well as details) have been captured. The raw range images captured ca. 1 GByte disc space. The resulting statue has ca. 5 millions of polygons and an accuracy of ca. 1 mm showing remarkable precision. The acquisition process last ca. 6 hours and the processing on a Sparc 10 one night. However, with several obvious optimizations the processing time can be accelerated to under one hour.

Other applications presented here include the modelling of buildings from different photographs.

Addresses

Michael Bender
Universität Kaiserslautern
Fachbereich Informatik
Postfach 3049
D-67618 Kaiserslautern
Germany
Tel. +49 631 205 4002
Fax. +49 631 205 3270
E-Mail: bender@informatik.uni-kl.de

Christian Bohn
GMD-IMK
Schloß Birlinghofen
D-53754 Sankt Augustin
Germany
Tel. +49 2241 14 2230
Fax. +49 2241 14 2040
E-Mail: bohn@gmd.de

Alan Chalmers
University of Bristol
Dept. of Computer Science
Merchant Ventures Building
Woodland Road
Bristol BS8 1UB
Great Britain
Tel. +44 117 954 5150
Fax. +44 117 954 5208
E-Mail: alan.chalmers@bristol.ac.uk

Michael Cohen
Microsoft Research
1 Microsoft Way
Redmond, WA 98052
USA
Tel. +1 206 703 0134
Fax. +1 206 936 0502
E-Mail: mcohen@microsoft.com

Philip Dutré
Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200 A
B-3030 Heverlee-Leuven
Belgium
Tel. +32 16 327 094
Fax. +32 16 327 996
E-Mail: Philip.Dutre@cs.kuleuven.ac.be

Dieter Fellner
Universität Bonn
Institut für Informatik III
Römerstraße 164
D-53117 Bonn
Germany
Tel. +49 228 550 315
Fax. +49 228 550 382
E-Mail: fellner@cs.uni-bonn.de

Alain Fournier
University of British Columbia
Computer Science
6356 Agricultural Rd
Vancouver BC V6T 1Z2
Canada
Tel.
Fax.
E-Mail: fournier@cs.ubc.ca

Bernd Fröhlich
Stanford University
CS Department
3B Gates Bldg
Stanford, CA 94305
USA
Tel. +1 415 725 5454
Fax. +1 415 723 0033
E-Mail: bernd@graphics.stanford.edu

Robert Garmann
Informatik VII
Universität Dortmund
44221 Dortmund
Germany
Tel. +49 231 755 6322
Fax. +49 231 755 6321
E-Mail:
garmann@ls7.informatik.uni-dortmund.de

Markus Groß
Institut für Informatik
ETH Zentrum
CH-8092 Zürich
Switzerland
Tel. +41 1 632 7114
Fax. +41 1 632 1172
E-Mail: gross@inf.ethz.ch

Hans Hagen
Universität Kaiserslautern
Fachbereich Informatik
Postfach 3049
D-67618 Kaiserslautern
Germany
Tel. +49 631 205 4071
Fax. +49 631 205 3270
E-Mail: hagen@informatik.uni-kl.de

Patrick M. Hanrahan
Stanford University
217 Center for Integrated Systems
Stanford, CA 94305-4070
USA
Tel. +1 415 723 8530
Fax.
E-Mail: hanrahen@cs.stanford.edu

Philip Jacob
Universität Kaiserslautern
Fachbereich Informatik
Postfach 3049
D-67618 Kaiserslautern
Germany
Tel. +49 631 205 2889
Fax. +49 631 205 3270
E-Mail: jacob@informatik.uni-kl.de

Frederik W. Jansen
Faculty of Technical Mathematics and Informatics
Delft University of Technology
P.O Box 356
NL-2600 Delft
The Netherlands
Tel. +31 15 278 5517
Fax. +31 15 278 7141
E-Mail: fwj@duticg.twi.tudelft.nl

Henrik Wann Jensen
Institute of Graphical Communication
Technical University of Denmark
Building 116
2800 Lyngby
Denmark
Tel. +45 45251650
Fax. +45 45938317
E-Mail: igkhwj@unidhp.uni-c.dk

Alexander Keller
Universität Kaiserslautern
Fachbereich Informatik
Postfach 3049
D-67618 Kaiserslautern
Germany
Tel. +49 631 205 3345
Fax. +49 631 205 3270
E-Mail: keller@informatik.uni-kl.de

Eric P. Lafortune
Department of Computer Science
Katholieke Universiteit Leuven
Celestijnenlaan 200A
B-3030 Heverlee
Belgium
Tel. +32 16 327094
Fax. +32 16 327996
E-Mail: eric.lafortune@cs.kuleuven.ac.be

Dani Lischinski
University of Washington
Dapt. of CSE
Box 352350
Seattle, WA 98195-2350
USA
Tel. +1 206 543 3368
Fax. +1 206 543 2969
E-Mail: danix@cs.washington.edu

Nelson Max
University of California
Lawrence Livermaore National Lab.
P.O.Box 808
Livermore CA 94550
USA
Tel. +1 510 422 4074
Fax. +1 510 423 8704
E-Mail: max@llnl.gov

Heinrich Müller
Universität Dortmund
Informatik VII
D-44221 Dortmund
Germany
Tel. +49 231 755 6324
Fax. +49 231 755 6321
E-Mail:
mueller@ls7.informatik.uni-dortmund.de

Stefan Müller
Fraunhofer Institut für
Graphische Datenverarbeitung
Wilhelminenstr. 7
D-64283 Darmstadt
Germany
Tel. +49 6151 155 188
Fax. +49 6151 155 199
E-Mail: stefanm@igd.fhg.de

Laszlo Neumann
Maros u. 36
H-1122 Budapest
Hungary
Tel. +36 1 554127
Fax. +36 1 3193985
E-Mail: neumann@odin.net

Gregory M. Nielson
Arizona State University
Department of Computer
Science and Engineering
Tempe, AZ 85287-5406
USA
Tel. +1 602 965 2785
Fax. +1 602 965 2751
E-Mail: nielson@asu.edu

Georg Pietrek
Universität Dortmund
Informatik VII
D-44221 Dortmund
Germany
Tel. +49 231 755 6122
Fax. +49 231 755 6321
E-Mail:
pietrek@ls7.informatik.uni-dortmund.de

Claude Puech
Universite Joseph Fourier
iMAGIS/GRAVIR-IMAG
B.P 53
F-38041 Grenoble Cedex 09
France
Tel. +33 76514690
Fax. +33 76446675
E-Mail: claude.puech@imag.fr

Xavier Pueyo
Universita de Girona
Dep. Informatica i Matematica Aplicada
Lluis Santaló s/n
E-17071 Girona
Spain
Tel. +34 72 418447
Fax. +34 72 418399
E-Mail: xavier@ima.udg.es

Werner Purgathofer
Institut für Computergraphik
TU Wien
Karlsplatz 13/186-2
A-1040 Wien
Austria

Tel. +43 1 58801 4548
Fax. +43 1 5874932
E-Mail: wp@cg.tuwien.ac.at

Georgios Sakas
Fraunhofer Institut für
Graphische Datenverarbeitung
Wilhelminenstraße 7
D-64283 Darmstadt
Germany
Tel. +49 6151 155 153
Fax. +49 6151 155 444
E-Mail: gsakas@igd.fhg.de

David Salesin
University of Washington
Dept. of Computer Science
and Engineering, FR-35
Seattle, WA 98195
USA
Tel. +1 206 685 1227
Fax. +1 206 543 2969
E-Mail: salesin@cs.washington.edu

Stephan Schäfer
Universität Bonn
Institut für Informatik III
Römerstraße 164
D-53117 Bonn
Germany
Tel. +49 228 550 391
Fax. +49 228 550 382
E-Mail: schaefer@graphics.cs.uni-bonn.de

Peter Schröder
California Institute of Technology
Computer Science Department 256-80
Room 286
Pasadena, CA 91125
USA
Tel. +1 818 395 4269
Fax. +1 818 792 4257
E-Mail: ps@cs.caltech.edu

Mark Segal
Silicon Graphics
2011 N. Shoreline Blvd.
P.O.Box 7311
Mountain View, CA 94039
USA
Tel. +1 415 933 1385
Fax. +1 415 933 0979
E-Mail: segal@sgi.com

Hans-Peter Seidel
Universität Erlangen
Graphische Datenverarbeitung
Am Weichselgarten 9
D-91058 Erlangen
Germany
Tel. +49 9131 85 9919
Fax. +49 9131 85 9931
E-Mail: seidel@informatik.uni-erlangen.de

Philipp Slusallek
Universität Erlangen
IMMD IX - GDV
Am Weichselgarten 9
D-91058 Erlangen
Germany
Tel. +49 9131 85 9925
Fax. +49 9131 85 9931
E-Mail: slusallek@informatik.uni-erlangen.de

Marc Stamminger
Universität Erlangen
IMMD IX - GDV
Am Weichselgarten 9
D-91058 Erlangen
Germany
Tel. +49 9131 85 9928
Fax. +49 9131 85 9931
E-Mail: stamminger@informatik.uni-erlangen.de

Wolfgang Stürzlinger
Johannes Kepler University
Dept. of Graphics and Parallel Processing
Altenbergstraße 69
A-4040 Linz
Austria
Tel. +43 732 2468 9228
Fax. +43 732 2468 9496
E-Mail: wrzl@gup.uni-linz.ac.at

Laszlo Szirmay-Kalos
Technical University of Budapest
Department of Process Control
Müegyetem rkp. 9. R. 216/5
H-1111 Budapest
Hungary
Tel. +36 1 463 2030
Fax. +36 1 463 2204
E-Mail: szirmay@fsz.bme.hu

Daniela Tost
Universitat Politecnica de Catalunya
Dept. Llenguatges i Sistemas Informatics
Edif. ETSEIB-UPC
Diagonal 647
E-08028 Barcelona
Spain
Tel. +34 3 401 6667
Fax. +34 3 401 6000
E-Mail: dani@ksi.upc.es