Report on Dagstuhl Seminar 0025

Image Synthesis and Interactive 3D Graphics

June 19 – 23, 2000

Organized by

Michael Cohen, Microsoft Research, Redmond, USA
Heinrich Müller, Universität Dortmund, Germany
Claude Puech, Université Joseph Fourier, Grenoble, France
Hans-Peter Seidel, Max-Planck Institut für Informatik, Saarbrücken, Germany

The goal of digital image synthesis and interactive 3D graphics is the calculation of images or image sequences from three-dimensional scenes. The task has many applications and challenging research problems which are particularly caused by the quest for rapid rendering from the requirements of applications on one side, and the inherently complex computational task on the other side. A typical problem is the development of approximating approaches which allow to calculate the optical effects sufficiently precise, but practically efficient. For the treatment of the problem, methods of finite element calculation, Monte-Carlo-Integration, design of efficient data structures and algorithms, but also human visual perception are applied. Another emphasis lies on the development of solutions on specialized hardware architectures and their use by adapted algorithms and software. In the last years, significant progress has been achieved, but by far not all important problems have been solved yet.

This Dagstuhl Seminar has provided a forum for leading researchers in this area to present their ideas. It has particularly benefited from the active participation of a high number of young researchers who have been for the first time at a Dagstuhl Seminar. Staying one week together at the beautiful place of Dagstuhl has stimulated the scientific and private exchange between the more than 50 participants from 11 countries far beyond that what can happen at usual conferences.

The number of participants and the willingness of active participation by giving a presentation caused 47 talks. By keeping the presentation time short there has been the necessary time for discussions which has been extensively and intensively used by the attendees. Although the presentations covered an extremely broad spectrum, it was surprising to always see a well-filled auditorium. The impression was that the chance was used to learn more on fields aside from the often very specialized own research. Hot topics, like e.g. point-based rendering, have found particular interest.

The positive feedback of the participants have shown again that there is a need of events of the type of Dagstuhl Seminars, and we have the hope to be able to have a follow-up in the future.

Table of Contents

Philipp Slusallek
Interactive Ray-Tracing

Per H. Christensen
Faster Photon Map Global Illumination

Henrik Wann Jensen
Global Illumination Test Scenes

Bernard Péroche
Global Illumination with Light Vectors

Marc Stamminger
Corrective Texturing

László Szirmay-Kalos
Unifying Random-Walk and Iteration in Global Illumination Rendering

Robert F. Tobler
The Advanced Rendering Toolkit - A Framework for Photorealistic Image Synthesis

Daniel Cohen-Or
Virtual Occluders and From Region Techniques

James Stewart
Cloth Shading

Phil Dutre
Visibility in Point Clouds

Rüdiger Westermann
Volume Rendering for the Consumer Market

Werner Purgathofer
Interactive Urban Visualisation

Philippe Bekaert
Monte Carlo Techniques for Radiosity

Alexander Keller
Strictly Deterministic Sampling Methods in Computer Graphics

Ingo Wald
An analysis of the photon map

Gregory M. Nielson
Progressive, Top-Down, Modeling of Volume Data Using Tetrahedral Coons Volumes

Heinrich Müller
Interactive manipulation of voxel volumes with free-formed voxel tools

Dani Tost
Visualization of Contour Surfaces

Ming C. Lin
HIGH-FIDELITY VIRTUAL TOUCH: Algorithms and Applications

Bernd Fröhlich
Geo-Scientific Visualization

Eduard Gröller
Interactive 3D Graphics in Medicine

Markus Gross
Surfels: An Alternative Rendering Primitive

Matthias Hopf
Accelerating Filters Using Graphics Hardware

Brian Curless
Surface Light Fields for 3D Photography

Dinesh Manocha
Interactive Walkthrough of Massive Models

Stefan Gumhold
Triangle Mesh Compression

Dani Lischinski
Streaming of Complex 3D Scenes for Remote Walkthroughs

Leif Kobbelt
$\sqrt{3}$-Subdivision

Gordon Müller
Hierarchical Data Structures for Efficient Rendering and Navigation

Reinhard Klein
3D Mesh-Editing

Real-time programmable shading for vertices and fragments

William R. Mark
Stanford University

Our research group has built a programmable-shading system that allows a single shader to include per-fragment, per-vertex, and per-primitive-group computations. The shading system runs in real-time on current graphics hardware. It currently uses the host CPU for per-vertex computations, but should be able to use programmable vertex hardware when it becomes available. Per-vertex computations are particularly useful for lighting, coordinate-space conversions, and texture-coordinate generation.

Our shading system supports separate surface and light shaders, in a manner similar to RenderMan. This design decision has system-wide implications, primarily because we allow a light shader to compute either a per-fragment or per-vertex result. Our system provides an immediate-mode geometry interface that hides the underlying multi-pass rendering implementation from the application.

SMASH: A Next-Generation API for Programmable Graphics Accelerators

Michael McCool
University of Waterloo

SMASH is a testbed for low-level graphics API concepts and is meant to be a concrete target for the development of extensions to OpenGL. Specifically, SMASH is syntactically and conceptually similar to OpenGL but supports, among other experimental features, a common programmable shader sub-API that is compatible with both multi-pass and single-pass implementations of shaders.

Arbitrary numbers of shader parameters of various types can be bound to vertices of geometric primitives using a simple immediate-mode mechanism. Run-time specification, manipulation, and compilation of shaders at various levels of resolution is supported, including integrated support for per-primitive, per-vertex and per-fragment shaders.

Implementation of real-time rendering effects using SMASH can be enhanced with metaprogramming toolkits and techniques, up to and including RenderMan-like shading languages. We give several implementations of a two-term separable BRDF approximation as examples.

Interactive Horizon Mapping

Peter-Pike J. Sloan, Michael F. Cohen
Microsoft Research

Shadows play an important role in perceiving the shape and texture of an object. While some previous interactive shadowing methods are appropriate for casting shadows on other geometry they can not be applied to bump maps (which contain no explicit geometry.) Horizon Mapping is a technique used to compute shadows for bump-mapped surfaces. We map the technique into modern graphics API's and extend it to account more accurately for the geometry of the underlying surface. We also use it to represent limited self-shadowing for pure geometry. In mapping the algorithm to hardware, we use a novel method to interpolate orientation in tangent space over the surface. We show results of self-shadowing at frame rates.

Illuminating Micro-Geometry based on Precomputed Visibility

Wolfgang Heidrich
MPI für Informatik

Many researchers have been arguing that geometry, bump maps, and BRDFs present a hierarchy of detail that should be exploited for efficient rendering purposes. In practice however, this is often not possible due to inconsistencies in the illumination for these different levels of detail. For example, while bump map rendering often only considers direct illumination and no shadows, geometry-based rendering and BRDFs will mostly also respect shadowing effects, and in many cases even indirect illumination caused by scattered light.

In this paper, we present an approach for overcoming these inconsistencies. We introduce an inexpensive method for consistently illuminating height fields and bump maps, as well as simulating BRDFs based on precomputed visibility information. With this information we can achieve a consistent illumination across the levels of detail.

The method we propose offers significant performance benefits over existing algorithms for computing the light scattering in height fields and for computing a sampled BRDF representation using a virtual gonioreflectometer. The performance can be further improved by utilizing graphics hardware, which then also allows for interactive display.

Finally, our method also approximates the changes in illumination when the height field, bump map, or BRDF is applied to a surface with a different curvature.

Interactive Environment Map Filtering

Jan Kautz
MPI für Informatik

Different methods have been proposed that use prefiltered environment maps to achieve glossy reflections. These methods have in common that a given static environment map has to be filtered according to a reflectance model, which is a time-consuming process and can only be done as a preprocessing step.

We propose to approximate a given BRDF in a specific way so that the corresponding filter kernel is radially symmetric and shift-invariant over the sphere. We then propose to use the parabolic map representation for the environment map, which allows the usage of the convolution operation supported by the graphics hardware to perform the filtering. The filtering process can then be done at interactive rates and does not need to be performed in a preprocessing step anymore.

Perception-based progressive global illumination solution

Karol Myszkowski
University of Aizu

A novel view-independent technique for progressive global illumination computations has been proposed that uses prediction of visible differences to improve both efficiency and effectiveness of physically-sound lighting solutions. The technique is a mixture of stochastic (density estimation) and deterministic (adaptive mesh refinement) algorithms that are used in a sequence optimized to reduce the differences between the intermediate and final images as perceived by the human observer in the course of lighting computations. The quantitative measurements of visibility were obtained using the model of human vision captured in the Visible Differences Predictor (VDP) developed by Daly. The VDP responses were used to support selection of the best component algorithms from a pool of global illumination solutions, and to enhance the selected algorithms for even better progressive refinement of the image quality. The VDP was also used to determine the optimal sequential order of component-algorithm execution, and to choose the points at which switch-over between algorithms should take place. The resulting technique was validated experimentally in terms of the accuracy of lighting simulation and the fidelity of images through a comparison with complex real-world environments (more details on the experiment can be found on the Web page: `http://www.u-aizu.ac.jp/labs/csel/atrium2/`).

Annette Scheel

Tone Mapping for Interactive Walkthroughs
MPI für Informatik

When a rendering algorithm has created a pixel array of radiance values the task of producing an image is not yet completed. In fact, to visualize the result the radiance values still have to be mapped to luminances, which can be reproduced by the used display. This step is performed with the help of tone mapping operators. These tools have mainly been applied to still images, but of course they are just as necessary for walkthrough applications, in which several images are created per second. In this talk it is shown how tone mapping can also be introduced into interactive radiosity viewers, where the tone mapping continuously adjusts to the current view of the user. The overall performance is decreased only moderately, still allowing walkthroughs of large scenes.

Time Dependent Tone Mapping

Frédo Durand
MIT Laboratory for Computer Science

Tone Mapping is crucial to depict the large range of luminance encountered in real scene with current display technology which still has rather low contrast. In this paper, we deal with the problem of *time-dependent tone mapping*, that is, the simulation of the time course of visual adaptation to provide a faithful lighting immersion for time-varying images. We propose a computational model of visual adaptation and the corresponding tone-mapping operator. Our model is based on psychophysical and physiological data and simulates the steady state as well as transient behavior.

Our simulation of light adaptation takes both time-varying multiplicative and subtractive mechanisms into account, as well as the static response compression. The process of slow dark adaptation is accounted for by simulating the chemical regeneration of photopigments and the *equivalent* light that they generate. We simulate chromatic adaptation, that is the ability to discount the color of the illuminant, its degree of completeness and its time course. We demonstrate our model for off-line tone-mapping and for interactive display of 3D scenes with large differences of luminance.

Thoughts on Global Illumination

Brian Smits
University of Utah

One hundred years ago Fredholm developed his theory of integral equations. Since then many fields have developed solution techniques for integral equations similar to the rendering equation. It seems unlikely that we can develop significantly better solution techniques than we currently have without looking at the aspects of our problem that make it different. Two of these differences are the perceptual aspects of our problem, and that we often desire a large sequence of images. The higher level perceptual aspects of our problem are poorly understood, both by us and by the perceptual psychologists, but seem likely to eventually provide the most significant benefit. Animation currently suffers due to using algorithms that were developed for static images.

Radiosity on Evolving Networks

Christian-Arved Bohn
GMD Institute for Media Communication

This work presents an alternative way for the discretization of a virtual environment for a radiosity finite element solution. Based on sample “rays” drawn from the geometry, a pair of *artificial neural networks* (ANN) is trained and their inner structures are interpreted as FEM meshing. Representations of the *light flow* as well as the *scene geometry* are generated. Based on this *virtual geometry* the required FEM integration operations are accomplished analytically. The network structure approaches the given geometry with the progress of *learning* and delivers a close functional description of the light flow in the virtual scene.

The main feature of this work is the new kind of generation of an FEM meshing for radiosity, which is completely independent from the scene objects. The neural network generates its own internal representation of the geometry, and thus, the approach is capable of detecting coherency or clusters of the underlying light flow in an optimal way. Due to the independence from the scene definition, *initial linking* is avoided and rough, fast solutions for even huge geometries are possible. A further essential feature is the combination of sampling and the FEM — importance sampling is driven by the light flow instead of examining the geometry relationships only.

Efficient lighting simulation for physiological plant growth

Cyril Soler
MPI für Informatik

Growing realistic virtual plants remains a challenge in terms of scientific modeling and computational algorithms. We demonstrate a complete simulator for plant growth, operating on physiological and biological factors, including the radiant energy reaching leaves of the plant. The distribution of illumination can be quite complex in large plants, because of shadowing, diffusion and transmission, and indirect lighting. We propose a new hierarchical instanciation algorithm for radiosity that enables an efficient, yet accurate determination of the illumination received by any part of the plant. In this algorithm, instances are equiped with suitable radiative properties and replace large amounts of geometry with controllable precision. This instanciation technique bridges the gap between explicit radiosity solutions and view-dependent approaches operating on instanced geometry. The algorithm also offers a promising approach to radiosity simulations in very complex architectural scenes involving some repetition. A complete lighting simulation algorithm for plants is presented, including hierarchical instanciation, the treatment of foliage translucency and heliotropism. Results s how that the instanciation technique allows the simulation of large trees that could not be dealt with previously, and permit to grow plants until a significant age, and observe the impact of the environment on their growth process.

Computer-Generated Pen-and-Ink Illustration of Trees

Oliver Deussen
University of Magdeburg, Germany

We present a method for automatically rendering pen-and-ink illustrations of trees. A given 3-d tree model is illustrated by the tree skeleton and a visual representation of the foliage using abstract drawing primitives. Depth discontinuities are used to determine what parts of the primitives are to be drawn; a hybrid pixel-based and analytical algorithm allows us to deal efficiently with the complex geometric data. Using the proposed method we are able to generate illustrations with different drawing styles and levels of abstraction. The illustrations generated are spatial coherent, enabling us to create animations of sketched environments. Applications of our results are found in architecture, animation and landscaping.

Illuminataion and Hierarchical Image-Based Rendering for Trees

Nelson Max
LLNL and UC Davis

This talk will cover three methods for illumination under trees. The first is for computing single-scattering illumination from the sun considered as a point source, by efficiently organizing the atmospheric illumiation information along slanted scan planes through the veiwpoint and the sun. The second produces the shadows under the trees from the sky illumation, with the sun included as an area rather than a point. These shadows are computed efficiently using fast Fourier transforms. The third involves computing multiple scattering of the light by reflection and transmission through the leaves, by treating the leaves as a volume distribution of infinitesimal planar fragments, whose areas and normals follow a distribution depending spatially only on the height above the ground.

In addition, I will describe hierarchical Image-Based Rendering methods for trees, first done in software, and more recently in hardware. The idea is to prepare precomputed views from the six + and - axis directions, including colors, normals, and depth at several layers, and then reproject them into a new viewing situation. These images are precomputed for all the different-sized objects in the hierarchical model of the tree, which are selected adaptively to give sufficient resolution for the current view. The visible surfaces are shaded using the color and (rotated) normal information in the reprojected view. By use of OpenGL features like texture mapping, alpha test, and color matrix transforms, the reprojection and shading can now be accomplised in hardware.

Lumigraph Rendering Through Warping

Hartmut Schirmacher
MPI für Informatik

We present a warping-based algorithm for the generation of arbitrary views from Lumigraphs at interactive frame rates. The algorithm partitions the Lumigraph's image plane in order to find the contribution of each input image to the final image. Only potentially contributing samples are actually reprojected. The final image is computed by depth-comparing and blending all samples that map to the same target pixel. Projection into the final view is factored into two steps, first the per-pixel warping within the original image plane, and second the final projection into the output view that is performed using OpenGL's texture mapping features.

The new method achieves interactive frame rates that are nearly independent of the number of input images, and makes use of both CPU and graphics hardware. Furthermore, the visual quality of the reconstructed views is superior to previous approaches, especially when only few input images are used.

Wavelet-based Light Field Compression

Ingmar Peter
Universität Tübingen

In the talk the compression of light field data using 4D nonstandard wavelet decomposition is desribed. For progressive transmission, storage, and rendering of the compressed light field data the new *wavelet stream* data structure is introduced. The wavelet coefficients are ordered in decreasing importance, encoding the position of the non discarded coefficients using significance maps.

Compared to the vector quantization compression method, the *wavelet stream* achieves up to 3 times higher compression ratios while producing the same error during reconstruction of the light field data. Despite of the high compression of the data, reasonable frame rates during light field rendering are achieved.

Automatic Image-Based Scene Model Acquisition and Visualization

Günther Greiner, Christian Vogelgsang, Benno Heigl
Universität Erlangen-Nürnberg

In the talk we present the design and realization of an integrated automatic framework for capturing, reconstructing and rendering real world scenes with image-based techniques. We use uncalibrated 2D video streams from a consumer video camera recording the scene data in an overlapping move pattern and reconstruct a calibrated view sequence from it. This view data is then compressed and transformed to a lightfield representation used for storage and rendering. The system design allows the usage of different compression techniques and lightfield parameterizations. We show the different steps of the process chain and describe the details of the underlying system architecture and its modular design which is capable of transparently supporting various different file formats and data representations required at different stages. As a result we present the application of the framework with a sample scene. Finally we outline current limitations and possibilities for future extensions of the system.

Creating Realistic 3-D Models from Many Camera Views

Eckehard Steinbach
Stanford University

In this talk we present an efficient method for calibrating many camera views simultaneously under the assumption that the silhouette of the object is vailable for all views. An error measure based on the mutual consistency of the object silhouettes in pairs of views is derived. We show how the partial derivatives of the error measure with respect to the calibration parameters can be computed analytically. The parameterization of our camera setup is assumed to be either unconstrained or constrained. In the former case, the position and orientation of each camera is independent of all other cameras. In the latter case, all cameras share some common calibration parameters leading to reduced dimensionality of the problem. Our experimental results suggest that the gradient-based minimization of the error measure robustly calibrates multiple camera views. Due to the explicit availability of gradients, the technique is also computationally efficient. Given the set of calibrated camera views we present a modified version of the voxel coloring algorithm to reconstruct a 3-D model of the object. Our voxel coloring strategy is based on a multi-hypothesis test of the voxel color in combination with explicit consideration of the finite extent of the voxel footprint observable in the image plane. This voxel coloring technique preserves fine detail of the object texture and leads to visually plasing 3D reconstructions.

Interactive Ray-Tracing

Philipp Slusallek
Saarland University

Ray-Tracing has some unique advantages over polygon rasterization such as implicit occlusion culling. This together with the long trend towards highly complex models has led many well-known researchers to predict that ray-tracing will eventually win over rasterization. Unfortunaty, this prediction has been wrong for the past 16 years.

However, there has been very little research that compares the two approaches on equal footing (such as the OpenGL geometry and shading capabilities) and tries to derive a sound basis for predicting the future. We approach these issues both from a software and a hardware perspective as we are exploring the space of high-performance ray-tracing architectures.

We have completed the development of a highly efficient scene traversal and ray-triangle intersection kernel on Intel P-III CPUs by exploiting their SIMD instruction set and optimizing for memory bandwidth and cache utilization. The test application renders common test scenes (Shirley/MGF) at 5-10 fps with simple shading on a single 600 MHz CPU, computing up to 1/2 million primary rays per second. Future work will include an IBR frontend, reordering of computations for even better cache utilization, handling of dynamic objects, distributed processing as well as continued work on designing a hardware version of the architecture. The software and hardware research will hopefully allow us to give a more precise answer on if and when ray-tracing might win over rasterization.

Faster Photon Map Global Illumination

Per H. Christensen
Square USA, Honolulu, Hawaii

The photon map method is an extension of ray tracing that enables it to efficiently compute caustics and soft indirect illumination on surfaces and in participating media. This presentation describes a method to further speed up the computation of soft indirect illumination on surfaces. The speed-up is based on the observation that the many look-ups in the global photon map during final gathering can be simplified by (pre)computing local irradiance values at the photon positions. Our tests indicate that the calculation of soft indirect illumination can be sped up by a factor of 5-7 in typical scenes. The irradiance values can be computed before or during rendering. Other applications include participating media and importance at "importon" positions.

Global Illumination Test Scenes

Henrik Wann Jensen
Stanford University

The global illumination community has discussed having a database of scenes that could be used to compare and validate different global illumination algorithms. We present a set of test scenes for global illumination algorithms and propose that they be the beginning of such a database. These scenes are designed to be as simple as possible and still test a particular aspect of energy transport. The scenes are available on a web site (http://www.cs.utah.edu/ bes/graphics/scenes/) in a variety of formats, along with images and pixel radiance data. We feel that the simplicity and availability of the models will make it easier for the community to use them, and that the field will benefit from a standard set of scenes.

Even though the test scenes are very simple they still contain modes of light transport that are not handled very well by most global illumination algorithms. For example one of the test scenes is made of three cubes arranged in such a way that all the visible light is indirect and even contains a caustics from a secondary light source. The fact that such a simple scene cannot be rendered efficiently with current global illumination algorithms show that there is still room for new research in this area.

Global Illumination with Light Vectors

Bernard Péroche
LISSE, Ecole des Mines de Saint Etienne

In order to try to solve the global illumination problem in a ray tracing environment, we introduced the concept of light vector. Inspired from the caching method developed by G. Ward and from the photon map introduced by H. Jensen, the light vector is suitable for any geometric model (not only polygons), for any material (the BRDF is taken into account) and works with area light sources and with participating media. For efficiency reasons, three kinds of light vectors are used: direct, indirect and caustici, and when some conditions are fulfilled, an interpolation of already computed light vectors allow to save computation time.

Corrective Texturing

Marc Stamminger
MPI für Informatik

We present a new hybrid rendering method for interactive walkthroughs in photometrically complex environments. The display process starts from some approximation of the scene rendered at high frame rates using graphics hardware. Additional computation power is used to correct this rendering towards a high quality ray tracing solution during the walkthrough. This is achieved by applying corrective textures to scene objects or entire object groups. These corrective textures contain a sampled representation of the differences between the hardware generated and the high quality solution. By reusing the textures, frame-to-frame coherence is exploited and explicit reprojections of point samples are avoided. Finally, we describe our implementation, which can display interactive walkthroughs of fairly complex scenes including high quality global illumination features.

Unifying Random-Walk and Iteration in Global Illumination Rendering

László Szirmay-Kalos
TU Budapest

The paper introduces a method that can combine continuous and finite-element approaches in a way, that the speed of finite-element based iteration and the accuracy of continuous random walks are preserved but without the excessive memory requirements of other finite-element approaches. The basic idea is to decompose the radiance function to a finite-element component that is only a rough estimate and to a difference component that is obtained by Monte-Carlo techniques. The classical iteration using finite-elements and random walks are handled uniformly in the framework of stochastic iteration. This uniform treatment allows the finite-element component to be built up adaptively aiming at minimizing the Monte-Carlo component. The method is also suited for interactive walkthoughs and view-animation since when the viewpoint changes only the small Monte-Carlo component needs to be recomputed. Using the approach quite complex scenes consisting of tens of thousands of surface elements can be rendered in about a minute, and when the solution is available, we can walk in the scene interactively.

The Advanced Rendering Toolkit - A Framework for Photorealistic Image Synthesis

Robert F. Tobler
VRVis Center for Virtial Reality and Visualization

We present a portable library for rendering photorealistic images which was designed with the followning main goals: scalability in scene database size and hardware performance, orthogonality of interfaces, and support of a variety of rendering algortihms.. The first goal was accomplished by using a scene graph with optimized accelerating structures. A number of orthogonal interfaces support the extension of the system with new geometric primitives, new surfaces (BRDFs), new Materials (phase-, albedo-, scattering-, and transfer funcions), new texture functions, new colour spaces, aso. The framework implements a few different rendering algorithms like raycasting, raytracing, path tracing, OpenGL z-buffer rendering and some global illumination pre-processing algorithms for raytracing.

Virtual Occluders and From Region Techniques

Daniel Cohen-Or
Tel Aviv University

In typical visibility culling algorithms the occlusion is tested from a point. Thus, these algorithms are applied in each frame during the interactive walkthrough. A more promising strategy is to find the Potentially visible set (PVS) from a region or viewcell, rather than from a point. The computation cost of the PVS from a viewcell would then be amortized over all the frames generated from the given viewcell.

Effective methods have been developed for indoor scenes, but for general arbitrary scenes, the computation of the visibility set from a region is more involved than from a point. Sampling the visibility from a number of view points within the region yields an approximated PVS, which may then cause unacceptable flickering artifacts during the walkthrough.

Conservative methods were introduced based on the occlusion of individual large convex objects. In these methods a given object or collection of objects is culled away if and only if they are fully occluded by a single convex occluder. However, it was shown that a convex occluder is effective only if it is larger than the viewcell.

Recently, new techniques were developed in which the visibility culling from a region is based on the combined occlusion of a collection of objects (occluder fusion). The collection or cluster of objects that contributes to the aggregate occlusion has to be neither connected nor convex. The effective from-region culling of these techniques is significantly larger than previous from-region visibility methods.

In my talk I'll introduce the notion of virtual occluders. Given a scene and a viewcell, a virtual occluder is a view-dependent (simple) convex object, which is guaranteed to be fully occluded from any given point within the viewcell and which serves as an effective occluder from the given viewcell. Virtual occluders compactly represent the aggregate occlusion for a given cell. The introduction of such view-dependent virtual occluders enables applying an effective region-to-region or cell-to-cell culling technique and efficiently computing a potential visibility set (PVS) from a region/cell. I'll present an object-space technique that synthesizes such virtual occluders by aggregating the visibility of a set of individual occluders. I'll show that only a small set of virtual occluders is required to compute the PVS efficiently on-the-fly during the real-time walkthrough.

Cloth Shading

James Stewart
University of Toronto

The goal of this research is to shade irregular surfaces, such as cloth. The principal problem here is to compute, for each point on the surface, how much of the outside environment is visible. It's from the outside environment that the direct illumination arrives, so the shading of each point can be estimated if the visible portion of the outside environment is known.

Algorithm Outline: Many planar slices of the surface are taken. In each slice, the planar 'visibility cone' from each surface point is computed with 2D visibility techniques. (The visibility cone from a point is the set of directions, in a given plane, in which the outside environment is visible.) Given that cone at a point, the direct primary illumination of the point can easily be computed for any distribution of illumination from the environment (but only from light travelling in that plane!). For a more accurate estimate of the visibility cones, and to remove the restriction of light travelling in the plane, several slices are taken at different orientations and sets of 2D cones are combined to form 3D cones.

Visibility in Point Clouds

Phil Dutre
Cornell University

In this report, we present a simple technique to evaluate the visibility between pairs of points in a scene. In most current rendering algorithms, visibility queries are evaluated exactly. Our approach approximates the visibility value between two points using a point cloud representation of the surfaces in the scene. The computed value is a function of the distance and orientation of points in the point cloud relative to the line segment connecting the two query points.

Volume Rendering for the Consumer Market

Rüdiger Westermann
Universität Stuttgart

Rendering techniques for 3D scalar field data have been developed to a high degree of sophistication. Among them, 3D texture based techniques have been proven to be superior to others with respect to rendering performance and ease of implementation. However, in order to bring volume rendering to the consumer market some problems still need to be addressed. In particular, bandwidth requirements have to be kept low and 2D textures have to be used until 3D textures become available.

In this talk we present a level-of-detail framework for 3D texture based volume rendering that allows for the rendering of large scale data sets even on low bandwidth graphics architectures. Furthermore, we outline new approaches for 2D texture based volume rendering using standard PC graphics adapters. In particular, we demonstrate how to take advantage of the multi-texture register combiners available on current Nvidia graphics processors in order to improve image quality and to considerably speed-up the rendering process.

Interactive Urban Visualisation

Werner Purgathofer, Dieter Schmalstieg, Michael Wimmer, Peter Wonka
TU Wien

Large city scenes have several properties that make interactive rendering simpler, despite the fact of consisting of huge data amounts. Buildings are non-changing objects that allow for preprocessing steps. Building walls are usually vertical forms without holes, that have one edge coincident with the ground. The viewer often moves around at constant height or at least with an upper bound to height. This talk presents how to extend cull maps to be used with view cells, so that for every place in a city a precomputed list of potentially visible building walls is available. This list is conservative in the sense that no actually visible wall can be lost, so that the pure rendering of these walls produces a correct image. However, walls are only representatives for complex geometry, including doors and windows. These can all be rendered at full geometric precision. With this technique it is possible to walk through a city model consisting of 8 million polygons in real time, i.e. the average frame time is below 1/60 of a second.

Monte Carlo Techniques for Radiosity

Philippe Bekaert
Katholieke Universiteit Leuven

The radiosity method is a physically based method to compute the illumination in a virtual environment with diffuse (matte) surfaces. It leads to very large systems of linear equations (100,000 equations is common). Moreover, the coefficients of these systems are given by non-trivial four-dimensional integrals (form factors). The Monte Carlo method is a method of last resort to solve such difficult mathematical problems. In this talk, an overview is presented of how the Monte Carlo method can be applied in the context of the radiosity problem.

The Monte Carlo method can be used for reliable form factor integration. It turns out to be difficult to determine how many samples are required for computing each form factor in such a strategy however.

The emphasis of the talk is therefore on Monte Carlo techniques for solving the radiosity system of linear equations directly, bypassing the need for explicit form factor computation by interpreting the form factors as probabilities that can be sampled efficiently by tracing rays. Such techniques include stochastic relaxation methods and random walk methods. Stochastic relaxation methods are stochastic adaptations of well known deterministic relaxation methods for linear systems: relaxation methods contain matrix-vector products, which can be estimated by Monte Carlo. Random walk methods for linear systems are very similar to random walk methods for integral equations, used in e.g. path tracing. They have been studied intensively since the second world war.

In order to overcome the slow convergence rate of Monte Carlo, various variance reduction techniques can be used. Variance reduction techniques for Monte Carlo radiosity include view-importance sampling, the use of a constant control variate, combinations of gathering and shooting and weighted importance sampling. Low-discrepancy sampling is another way to speed up Monte Carlo radiosity.

In order to reduce discretisation artifacts, various ways to incorporate hierarchical refinement and to compute higher order radiosity approximations have been proposed in the context of Monte Carlo radiosity.

In my experience, the resulting Monte Carlo radiosity algorithms are an alternative of great value for deterministic radiosity algorithms: they often yield useful images more rapidly, they require less storage and they are more reliable and easier to implement correctly and to use.

Among all discussed algorithms, in particular the simple stochastic Jacobi iterative method proposed by Neumann et al is a good candidate to use in practice. The basic algorithm is as good as other algorithms, but variance reduction techniques appear to be more efficient and easier to implement.

The talk is concluded with some directions for further improvement of Monte Carlo radiosity algorithms.

Strictly Deterministic Sampling Methods in Computer Graphics

Alexander Keller
University of Kaiserslautern

We introduce a strictly deterministic, meaning non-random, rendering method, which performs superior to state of the art Monte Carlo techniques. Its simple and elegant implementation on parallel computer architectures is capable of simulating anti-aliasing, motion blur, depth of field, area light sources, glossy reflection and transmission, participating media, and global illumination. We also illustrate how the design of quasi-Monte Carlo algorithms, i.e. strictly deterministic sampling methods based on number theory, is related to Monte Carlo algorithms based on probability theory.

An analysis of the photon map

Ingo Wald
Saarland University

The photon map is a very powerful method for efficiently illuminating complex environments. However, violating the underlying reconstruction assumptions can lead to visible artifacts. After analyzing the reasons for these reconstruction artifacts, we present new techniques for improving the photon map, namely a faster two-level approach for calculating the radiance estimate, the use of multiple photon maps and two importance sampling schemes for importance driven generation of photon maps and for faster direct lighting calculation.

Progressive, Top-Down, Modeling of Volume Data Using Tetrahedral Coons Volumes

Gregory M. Nielson
Arizona State University

A new method for modeling and compressing volume data is described. The method is adaptive and based upon tetrahedral decomposition of the domain. The adaptive features are more readily accommodated by not requiring a valid tetrahedral decomposition. A model based upon piece wise Coons volumes rather than piece wise linear elements leads to continuous models over this non conforming cellular domain. We illustrate the power and versatility of this new scheme applied to several three dimensional volume data sets.

Interactive manipulation of voxel volumes with free-formed voxel tools

Heinrich Müller, Jörg Ayasse
University of Dortmund

We are given a voxel workpiece and a voxel tool. The tool moves along a path with six degrees of freedom. The task is to update the voxels of the workpiece which are touched by the tool on its path, and display the results. Updating means that an affected workpiece voxel gets or looses material, depending on the tool mode which can be "applying" or "removing". The particular challenge is to reach the speed required for interactive manipulation. Our solution reaches this goal by reducing the set of all tool voxels to a small set of affecting voxels. The affecting voxels are determined in a preprocessing step for twelve elementary motions. From those, the affecting tool voxels of an arbitrary motion are obtained as the union over its components of elementary motion. Another issue of gaining speed is to decompose the given motion into "small" motions. For every small motion, the motion of an affecting voxel is replaced with a hull whose intersection with the workpiece determines the affected voxels. That approach causes a slight error which, however, is not too severe for many applications. Visualization is efficiently performed by an incremental marching cubes modification. On a SGI Octane MXE, interactive speed is reached for workpieces up to 400^3 voxels and tools up to 80^3 voxels. The algorithm is part of our Grid-Based Modeler (GBM) for two-handed interactive modeling in a CAVE-like environment.

Visualization of Contour Surfaces

Dani Tost, Anna Puig
Polytechnical University of Catalunya

Contour surfaces are composed of sets of plannar contours. They are the natural output of surface extraction algorithms based on contouring features in parallel image slices of volume models. They are also suitable for the representation of CAD objects with tubular elongated shapes such as pipes and tools. Rendering these surfaces consists in tiling between successive contours. Tiling is often performed in a preprocess that constructs auxiliary triangle mesh models which are generally large and that are only useful for visualization purposes.

We analyze means of representing the surfaces such that tiling can be performed on the fly during the rendering stage. This avoids the need of a double representation and it allows adaptive levels of resolution in the rendering. Surfaces are represented as graphs with generalized cylinders at the edges and branching information at the nodes. The generalized cylinders are defined by a skeleton curve, sets of contours and, eventually, a sweeping cross section. Contours can be represented parametrically or as an explicit set of points. Parametric contours are stored at points where the skeleton curve has a significant gradient or at points where the radius derivative changes so that additional cross sections can thus be interpolated during tiling. Tiling between these sections consists in clipping the sweeping cross section, if any, against the circular section or aproximating the circle with a polygon. Non-circular sections are represented explicitely as sets of adjacents curve segments composed of points expressed in polar coordinates. Tiling between equivalent segments is performed according to a one-to-one equivalence between points inside the segments and triangle fans are constructed at the segments ends to equilibrate the number of points per segment. The visualization is fast and easily tunable. It has been applied to the rendering of the cerebral vascular tree and the brain and it is being used in a tool machining simulation application.

HIGH-FIDELITY VIRTUAL TOUCH: Algorithms and Applications

Ming C. Lin
University of North Carolina at Chapel Hill

Extending the frontier of visual computing, haptic interfaces have the potential to increase the quality of human-computer communication by accommodating the sense of touch. They also provide an attractive augmentation to visual and aural display, and enhance the level of immersion in a virtual world. In this talk, I'll describe some recent research that we have been pursuing and the results that we have been able to achieve so far.

- H-COLLIDE: a framework for fast and accurate collision detection for haptic interaction with polygonal models
- inTouch: interactive multiresolution modeling and 3D painting with a haptic interface
- 6-DOF Haptic Rendering: force/torque display and applications

To conclude, I will briefly describe our work in progress and future research directions on 6-DOF haptic rendering, as well as some possible applications of haptics research.

Geo-Scientific Visualization

Bernd Fröhlich
GMD Institute for Media Communication

In this talk we describe tools and techniques for the exploration of geo-scientific data from the oil and gas domain in stereoscopic virtual environments. The two main sources of data in the exploration task are seismic volumes and multivariate well logs of physical properties down a bore hole. We have developed a props-based interaction device called the cubic mouse to allow more direct and intuitive interaction with a cubic seismic volume. The device consists of a cube-shaped box with three perpendicular rods passing through the center and buttons on the top for additional control. The rods represent the X, Y, and Z axes of a given coordinate system. Pushing and pulling the rods specifies constrained motion along the corresponding axes. Twisting the rods typically rotations around the corresponding axes. Embedded within the device is a six degree of freedom tracking sensor, which allows the rods to be continually aligned with a coordinate system located in a virtual world. This device effectively places the seismic cube in the user's hand. We have also integrated the device with other visualization systems for crash engineers and flow simulations. In these systems the Cubic Mouse controls the position and orientation of a virtual model and the rods move three orthogonal cutting or slicing planes through the model.

We have also developed a 3D texture based multi-resolution approach for handling massive volumetric data sets predominant in the oil and gas industry. Due to the restricted texture memory available and the limited bandwidth into the texture memory, these data sets cannot be rendered at full resolution. Our approach uses a hierarchical paging technique to guarantee a given frame rate. This technique displays lower resolutions of the data when a slice or volume rendering is moved fast through the data set, and fills in the high resolution, when the user slows down or stops. This behaviour correlates really well with motion blur.

Interactive 3D Graphics in Medicine

Eduard Gröller
TU Wien

A fast virtual endoscopy system for transbronchial biopsy is discussed. A cube is moved along a precalculated central path. Layers of environment maps (surface rendering, direct volume rendering) are precalculated. During navigation (move, zoom, pan) these layers are blended and displayed according to the current camera parameters. The web-based system is intended for training and operation planning purposes. The second system discussed is J-Vision. It is a Java front-end to a PACS system. In addition to 2D oriented data display (slices, multiplanar reconstruction, maximum intensity projection) 3D iso-surfaces are included in the functionality. A variation of the shear-warp algorithm ensures interactive display and modification of several iso-surfaces even on hardware without advanced graphics hardware support.

Surfels: An Alternative Rendering Primitive

Markus Gross
ETH Zürich

Surface elements (surfels) are a powerful paradigm to efficiently render complex geometric objects at interactive frame rates. Unlike classical surface discretizations, i.e., triangles or quadrilateral meshes, surfels are point primitives without explicit connectivity. Surfel attributes comprise depth, texture color, normal, and others. As a pre-process, an octree-based surfel representation of a geometric object is computed. During sampling, surfel positions and normals are optionally perturbed, and different levels of texture colors are prefiltered and stored per surfel. During rendering, a hierarchical forward warping algorithm projects surfels to a z-buffer. A novel method called visibility splatting determines visible surfels and holes in the z-buffer. Visible surfels are shaded using texture filtering, Phong illumination, and environment mapping using per-surfel normals. Several methods of image reconstruction, including supersampling, offer flexible speed-quality tradeoffs. Due to the simplicity of the operations, the surfel rendering pipeline is amenable for hardware implementation. Surfel objects offer complex shape, low rendering cost and high image quality, which makes them specifically suited for low-cost, real-time graphics, such as games.

Accelerating Filters Using Graphics Hardware

Matthias Hopf
Universität Stuttgart

Filtering is an important part in scientific visualization. It is used for noise reduction, feature extraction, segmentation, and registration. Multiresolution techniques also generally rely on filters for their operation. In order to accelerate the typical visualization cycle, we use standard OpenGL graphics hardware to implement certain common filtering techniques, avoiding the necessary texture reload and exploiting the typically high memory bandwidth of modern graphics cards.

By this approach three dimensional linear filters, morphological operators, and wavelet analysis can be performed at three to six times the speed of fine tuned software implementations.

Additionaly we have implemented a benchmark system, distributed under GPL, which is capable of running basic OpenGL commands with arbitrary parameters that can be specified on the command line. With that system we are able to benchmark algorithms on new platforms, where the original code has not been ported to, with an estimated error bound of 1algorithms, as long as they decompose nicely into a set of basic OpenGL commands with typical parameter sizes.

Surface Light Fields for 3D Photography

Brian Curless
University of Washington

A surface light field is a function that assigns a color to each ray originating on a surface. Surface light fields are well suited to constructing virtual images of shiny objects under complex lighting conditions. This paper presents a framework for construction, compression, interactive rendering, and rudimentary editing of surface light fields of real objects. Generalizations of vector quantization and principal component analysis are used to construct a compressed representation of an object's surface light field from photographs and range scans. A new rendering algorithm achieves interactive rendering of images from the compressed representation, incorporating view-dependent geometric level-of-detail control. The surface light field representation can also be directly edited to yield plausible surface light fields for small changes in surface geometry and reflectance properties.

Interactive Walkthrough of Massive Models

Dinesh Manocha
University of North Carolina at Chapel Hill

Computer-aided design (CAD) applications and scientific visualizations often need user-steered interactive displays or walkthroughs of very complex environments. Structural and mechanical designers often create models of ships, oil platforms, spacecraft, process plants and urban environments, composed of tens of millions of primitives, whose complexity exceeds the interactive visualization capabilities of current graphics systems. Yet for such structures the design process, and especially the multi-disciplinary design review process, benefits greatly from interactive walkthroughs.

In this talk we give an overview of our recent work on interactive walkthrough of massive models. These include new algorithmic approaches for rendering acceleration based on visibility culling, model simplification and image-based representations. We will initially new and present general algorithms for model simplification. They produce high quality and drastic approximations and we also compute tight error bounds. In the second part of the talk, we present a framework for rendering very large models at nearly interactive rates. The framework scales with model size. Our framework can integrate multiple-rendering acceleration techniques, including visibility culling, geometric levels-of-detail and image-based approaches. We describe the database representation scheme for massive models used by the framework and a pipeline to manage the allocation of system resources among different techniques. We demonstrate the system on a model of a coal-fired power plant composed of more than 15 million triangles and a DoubleEagle Tanker composed of more than 82 million triangles.

Triangle Mesh Compression

Stefan Gumhold
Universität Tübingen

In this talk single resolution triangle mesh compression techniques are compared according to their algorithmic complexity and their encoding efficiency. The coding efficiencies of the best methods for single resolution, progressive and non manifold progressive triangle mesh compression are compared. In the main part the compression method of the Cut-Border Machine is presented and the coding efficiency is improved by exploiting more connectivity information of the so far compressed or decompressed part. Measurements show that the Cut-Border Machine is not only very fast, but also allows to encode typical triangle meshes with two bits per vertex. The last section of the talk discusses current work of other authors on polygonal mesh compression. In particular the Face Fixer approach is presented.

Streaming of Complex 3D Scenes for Remote Walkthroughs

Eyal Teler, Dani Lischinski
The Hebrew University of Jerusalem

We describe a 3D scene streaming approach for remote walkthroughs. In a remote walkthrough, a user on a client machine interactively navigates through a scene that resides on a remote server. Our approach allows a user to walk through a remote 3D scene, without having to download the entire scene from the server. Our algorithm achieves this by selectively transmitting only small parts of the scene, and lower quality representations of objects, based on the user's viewing parameters and the available connection bandwidth. An online optimization algorithm selects which object representations to send, based on the integral of a benefit measure along the predicted path of movement. Rendering quality depends on the available bandwidth, but practical navigation of the scene is possible even when bandwidth is low.

$\sqrt{3}$-Subdivision

Leif Kobbelt
MPI für Informatik

A new stationary subdivision scheme is presented which performs *slower* topological refinement than the usual dyadic split operation. The number of triangles increases in every step by a factor of 3 instead of 4. Applying the subdivision operator *twice* causes a uniform refinement with *tri*-section of every original edge (hence the name $\sqrt{3}$-subdivision) while two dyadic splits would *quad*-sect every original edge. Besides the finer gradation of the hierarchy levels, the new scheme has several important properties: The stencils for the subdivision rules have minimum size and maximum symmetry. The smoothness of the limit surface is C^2 everywhere except for the extraordinary points where it is C^1. The convergence analysis of the scheme is presented based on a new general technique which also applies to the analysis of other subdivision schemes. The new splitting operation enables locally adaptive refinement under built-in preservation of the mesh consistency without temporary crack-fixing between neighboring faces from different refinement levels. The size of the surrounding mesh area which is affected by selective refinement is smaller than for the dyadic split operation. We further present a simple extension of the new subdivision scheme which makes it applicable to meshes with boundary and allows us to generate sharp feature lines.

Hierarchical Data Structures for Efficient Rendering and Navigation

Gordon Müller
TU Braunschweig

The interactive exploration of highly complex 3D scenes containing hundreds of thousands of objects is one of the major challenges in computer graphics. In this talk we will present an algorithm for the creation of object hierarchies that are well suited to accelerate the visualization of complex scenes. Based on the object hierarchy many non-visible objects are detected by testing inner nodes of the scene graph for conservative visibility (occlusion culling) based on bounding box projection. Furthermore a technique is presented that eliminates many superficious visibility tests by exploiding frame-to-frame coherency in interactive applications.

3D Mesh-Editing

Reinhard Klein
TU Darmstadt

We present a new approach to the editing of triangulated surface meshes. The main features of our new algorithm are that we can exactly specify the area that is being edited, the area that remains unchanged and an area that serves as a stiff handle which is moved but doesn't change its shape. This allows intuitive editing of arbitrary triangle meshes. Furthermore, the shape of the deformed area can be modified interactively by manipulating a simple parametric curve, e.g. a Bspline. There are no special requirements on the connectivity of the triangle meshes. In spite of these features the algorithm is very fast and simple and can easily be generalized to the case of tetrahedrized volumes.

Addresses

Mark Alexa
Technische Universität Darmstadt
FB Graphische Datenverarbeitung
Rundeturmstr. 6
D-64283 Darmstadt
Germany
Tel: +49-6151-155-674
Fax: +49-6151-155-199
e-Mail: alexa@gris.informatik.tu-darmstadt.de

Philippe Bekaert
Katholieke Universiteit Leuven
Dept. of Computer Science
Celestijnenlaan 200 A
B-3001 Heverlee - Leuven
Belgium
Tel: +32-16-32.76.67
Fax: +32-16-32.79.96
e-Mail: Philippe.Bekaert@cs.kuleuven.ac.be
URL: www.cs.kuleuven.ac.be/˜philippe/

Christian Bohn
GMD
IMK - Institute for Media Communication
Schloss Birlinghoven
D-53754 St. Augustin
Germany
e-Mail: bohn@gmd.de
URL: imk.gmd.de/reta

Per Christensen
Square USA Inc.
Honolulu Studio
Suite 3100
55 Merchant Street
HI 96813 Honolulu
USA
Tel: +1-808-545-4861
e-Mail: per.christensen@acm.org
URL: www.aloha.net/˜myandper/

Daniel Cohen-Or
Tel Aviv University
School of Mathematical Sciences
Ramat-Aviv
69978 Tel-Aviv
Israel
Tel: +972-3-640-5368
Fax: +972-3-640-9357
e-Mail: daniel@math.tau.ac.il
URL: www.math.tau.ac.il/˜daniel

Brian Curless
University of Washington
Dept. of Computer Science & Engineering
114 Sieg Hall
Box 352350
WA 98195-2350 Seattle
USA
Tel: +1-206 685-3796
Fax: +1-206 543-2969
e-Mail: curless@cs.washington.edu
URL: www.cs.washington.edu/homes/curless

Oliver Deussen
Universität Magdeburg
Institut für Simulation u. Graphik
Postfach 41 20
D-39016 Magdeburg
Germany
Tel: +49-391-671-1332
Fax: +49-391-671-1164
e-Mail: deussen@acm.org
URL: wwwisg.cs.uni-magdeburg.de/˜deussen/

Frédo Durand
MIT
Laboratory for Computer Science
NE43-255 / Computer Graphics Group
200 Technology Square
MA 02139 Cambridge
USA
Fax: +1-617-253-4640
e-Mail: fredo@graphics.lcs.mit.edu
URL: graphics.lcs.mit.edu/˜fredo/

Philip Dutre
Cornell University
Program of Computer Graphics
580 Rhodes Hall
NY 14853-3801 Ithaca
USA
Tel: +1-607 255-2067
Fax: +1-607-255-0806
e-Mail: phil@graphics.cornell.edu
URL: www.graphics.cornell.edu/˜phil/

Thomas Ertl
Universität Stuttgart
Institut für Informatik
Breitwiesenstraße 20-22
D-70565 Stuttgart
Germany
Tel: +49-711-7816-331
Fax: +49-711-7816-340
e-Mail: ertl@informatik.uni-stuttgart.de
URL: wwwvis.informatik.uni-stuttgart.de/

Dieter Fellner
TU Braunschweig

Institut für Computer Graphik
Rebenring 18
D-38106 Braunschweig
Germany
Tel: +49-531-391-2101
Fax: +49-531-391-2103
e-Mail: d.fellner@tu-bs.de
URL: www.cg.cs.Tu-Bs.de/people/Fellner

Bernd Fröhlich
GMD
IMK - Institute for Media Communication
Schloss Birlinghoven
D-53754 St. Augustin
Germany
e-Mail: bernd.froehlich@gmd.de
URL: imk.gmd.de/docs/ww/people.mhtml

Günther Greiner
Universität Erlangen-Nürnberg
Informatik 9
Am Weichselgarten 9
D-91058 Erlangen
Germany
Tel: +49-9131-852-9919
Fax: +49-9131-852-9931
e-Mail: greiner@informatik.uni-erlangen.de
URL: www9.informatik.uni-erlangen.de/

Markus Gross
ETH Zürich
Dept. Informatik
CH-8092 Zürich
Switzerland
Tel: +41-1-632-7114
Fax: +41-1-632-1172
e-Mail: grossm@inf.ethz.ch
URL: www.graphics.ethz.ch

Eduard Gröller
Technische Universität Wien
Inst. für Computergraphik
Karlsplatz 13 / 1862
A-1040 Wien
Austria
Tel: +43-1-58801-18682
Fax: +43-1-58801-18698
e-Mail: groeller@cg.tuwien.ac.at
URL: www.cg.tuwien.ac.at/home/

Stefan Gumhold
Universität Tübingen
Wilhelm-Schickard-Institut für Informatik, GRIS
C9 P19
Auf der Morgenstelle 10, C9
D-72076 Tübingen
Germany
Fax: +49-7071-295-466
e-Mail: sgumhold@gris.uni-tuebingen.de
URL: www.gris.uni-tuebingen.de/˜sgumhold/

Wolfgang Heidrich
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: 0681-9325-429
Fax: 0681-9325-499
e-Mail: heidrich@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/˜heidrich/

Thorsten Holtkämper
GMD
Institut fur Methodische Grundlagen
Schloss Birlinghoven
Postfach 13 16
D-53731 St. Augustin
Germany
Tel: +49 / 2241-14 2103
Fax: +49 / 2241-14 2040
e-Mail: holtkaemper@gmd.de
URL: viswiz.gmd.de/˜tholt/

Matthias Hopf
Universität Stuttgart
Institut für Informatik
Breitwiesenstraße 20-22
D-70565 Stuttgart
Germany
Tel: +49-711-7816-404
Fax: +49-711-7816-340
e-Mail: hopf@informatik.uni-stuttgart.de
URL: wwwvis.informatik.uni-stuttgart.de/˜hopf/

Henrik Wann Jensen
Stanford University
Dept. of Computer Science
Computer Graphics Lab. - Gates Bldg. - Room 362B
CA 94305-9040 Stanford
USA
Tel: +1-650-725-3696
Fax: +1-650-723-0033
e-Mail: henrik@graphics.stanford.edu
URL: www.graphics.stanford.edu/˜henrik

Jan Kautz
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: +49-681-9325-427
Fax: +49-681-9325-499
e-Mail: kautz@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/˜jnkautz

Alexander Keller
Universität Kaiserslautern
FB Informatik
36/212
Postfach 3049
D-67653 Kaiserslautern
Germany
Tel: +49-631-205 3345
Fax: +49-631-205 3270
e-Mail: keller@informatik.uni-kl.de
URL: www.uni-kl.de/AG-Heinrich/Alex.html

Reinhard Klein
Technische Universität Darmstadt
FB Graphische Datenverarbeitung
Rundeturmstr. 6
D-64283 Darmstadt
Germany
Tel: +49-7071-29 54 62
Fax: +49-7071-29 54 66
e-Mail: reinhard@gris.uni-tuebingen.de
URL: www.gris.uni-tuebingen.de/people/staff/reinhard.html

Leif Kobbelt
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: ++49 (681) 9325-408
Fax: ++49 (681) 9325-499
e-Mail: kobbelt@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/˜kobbelt/

Eric Lafortune
Luciad
Parijsstraat 74
B-3000 Leuven
Belgium
Tel: +32-16-23 95 91
Fax: +32-16-29 34 22
e-Mail: eric.lafortune@luciad.com
URL: www.graphics.cornell.edu/˜eric/

Ming Lin
University of North Carolina at Chapel Hill
Dept. of Computer Science
CB #3175
Sitterson Hall
NC 27599-3175 Chapel Hill
USA
Tel: +49-(919) 962-1974
Fax: +49-(919) 962-1799
e-Mail: lin@cs.unc.edu
URL: www.cs.unc.edu/˜lin/

Dani Lischinski
The Hebrew University of Jerusalem
School of Computer Schience & Engineering
Givat Ram
91904 Jerusalem
Israel
Tel: +972-2-658-5087
Fax: +972-2-658-5439
e-Mail: danix@cs.huji.ac.il
URL: www.cs.huji.ac.il/˜danix

Dinesh Manocha
University of North Carolina at Chapel Hill
Dept. of Computer Science
CB #3175
Sitterson Hall
NC 27599-3175 Chapel Hill
USA
Tel: +1-919-962 1749
Fax: +1-919-962 1799
e-Mail: manocha@cs.unc.edu
URL: www.cs.unc.edu/˜dm/

Bill Mark
Stanford University
Graphics Laboratory
Rm 364-3B
Gates Building
CA 94305 Stanford
USA
Fax: +1-650-723-0033
e-Mail: billmark@graphics.stanford.edu
URL: graphics.stanford.edu/˜billmark/

Nelson Max
University of California
Lawrence Livermore National Lab.
L-560
7000 East Ave.
CA 94551 Livermore
USA
Tel: +1-925-422-40 74
Fax: +1-925-422-6287
e-Mail: max2@llnl.gov
URL: www.llnl.gov/graphics/

Michael McCool
University of Waterloo
Dept. of Computer Science
200 University Avenue West
ON-N2L 3G1 Waterloo
Canada
Tel: +1-(519) 888-4567 x4422
Fax: +1-(519) 885-1208
e-Mail: mmccool@cgl.uwaterloo.ca
URL: www.cgl.uwaterloo.ca/˜mmccool

Gordon Müller
TU Braunschweig
Institut für ComputerGraphik
Rebenring 18
D-38106 Braunschweig
Germany
Tel: +49-531-391-2104
Fax: +49-531-391-2103
e-Mail: gordon.mueller@tu-bs.de
URL: www.cg.cs.tu-bs.de/people/mueller

Heinrich Müller
Universität Dortmund
FB Informatik VII
D-44221 Dortmund
Germany
Tel: +49-231-755-63 24
Fax: +49-231-755-63 21
e-Mail: mueller@ls7.informatik.uni-dortmund.de
URL: ls7-www.informatik.uni-dortmund.de/

Karol Myszkowski
University of Aizu
Software Dept.
Tsuruga Ikki-machi
965-8580 Aizu Wakamatsu
Japan
Fax: +81-242-372-747
e-Mail: k-myszk@u-aizu.ac.jp
URL: www.u-aizu.ac.jp/˜k-myszk/

László Neumann
Maros u. 36
H-1122 Budapest
Hungary
Tel: +36-1-203-7310
e-Mail: lneumann@cg.tuwien.ac.at

Gregory M. Nielson
Arizona State University
Dept. of Computer Science and Engineering
Box 87 54 06
AZ 85287-5406 Tempe
USA
Tel: +1-602-965-27 85
Fax: +1-602-965-27 51
e-Mail: nielson@asu.edu
URL: www.eas.asa.edu

Bernard Peroche
Ecole des Mines de St. Etienne
Dept. LISSE
Cours Fauriel
F-42023 St. Etienne
France
Tel: +33-477 420168
Fax: +33-477 426666
e-Mail: peroche@emse.fr

Ingmar Peter
Universität Tübingen
Wilhelm-Schickard-Institut für Informatik, GRIS
Auf der Morgenstelle 10, C9
D-72076 Tübingen
Germany
Tel: 07071-29-75464
Fax: 07071-29-5466
e-Mail: peter@gris.uni-tuebingen.de
URL: www.gris.uni-tuebingen.de/˜peter

Claude Puech
Université Joseph Fourier
Laboratoire IMAG
F-38041 Grenoble
France
Tel: +33-4-76.51.46.90
Fax: +33-4-76.44.66.75
e-Mail: claude.puech@imag.fr

Xavier Pueyo
Universitat de Girona
Institut d'Informatica i Aplicacions
Campus de Montilivi
E-17071 Girona
Spain
Tel: +34-72-418-447
Fax: +34-72-418-792
e-Mail: xavier.pueyo@ima.udg.es

Werner Purgathofer
Technische Universität Wien
Inst. für Computergraphik
Favoritenstr. 9-11
A-1040 Wien
Austria
Tel: +43-1-58801-18648 (Sekr.18602)
Fax: +43-1-58801-18698
e-Mail: wp@cg.tuwien.ac.at
URL: www.cg.tuwien.ac.at

Annette Scheel
MPI für Informatik
AG 4
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: +49-681-9325-420
Fax: +49-681-9325-499
e-Mail: scheel@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/

Hartmut Schirmacher
MPI für Informatik

Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Fax: +49-681-9325-499
e-Mail: schirmacher@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/˜htschirm/

Hans-Peter Seidel
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: +49-681-9325-400
Fax: +49-681-9325-499
e-Mail: hpseidel@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/units/ag4/people.html

Peter-Pike Sloan
Microsoft Corporation
1 Microsoft Way
WA 98052-6399 Redmond
USA
Tel: +1-425-936-4246
e-Mail: ppsloan@microsoft.com

Philipp Slusallek
Universität des Saarlandes
FR Informatik
Im Stadtwald
PF 15 11 50
D-66041 Saarbrücken
Germany
Tel: +49-681-302-3830
Fax: +49-681-302-3843
e-Mail: slusallek@cs.uni-sb.de
URL: graphics.cs.uni-sb.de/˜slusallek

Brian Smits
University of Utah
Computer Science Dept.
Room 3190
50 South Central Campus Drive
UT 84112 Salt Lake City
USA
Tel: +49-801-585-5355
Fax: +49-801-581-5843
e-Mail: bes@cs.utah.edu
URL: www2.cs.utah.edu/˜bes/

Cyril Soler
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: (+49) 681 9325 426
Fax: (+49) 681 9325 499
e-Mail: soler@mpi-sb.mpg.de
URL: www-imagis.imag.fr/Membres/Cyril.Soler/

Marc Stamminger
MPI für Informatik
Stuhlsatzenhausenweg 85
D-66123 Saarbrücken
Germany
Tel: +49 681 9325 422
Fax: +49 681 9325 499
e-Mail: stamminger@mpi-sb.mpg.de
URL: www.mpi-sb.mpg.de/˜mcstammi/

Eckehard Steinbach
Stanford University
Information Systems Laboratory
350 Serra Mall
CA 94305 Stanford
USA
Tel: +1-650-724-3648
e-Mail: steinb@stanford.edu
URL: www.stanford.edu/˜steinb

James Stewart
University of Toronto
Dept. of Computer Science
10 King's College Road
ON-M5S 3G4 Toronto
Canada
Tel: +416-978-5359
Fax: +416-978-5184
e-Mail: jstewart@cs.toronto.edu
URL: www.dgp.toronto.edu/˜jstewart

László Szirmay-Kalos
Technical University of Budapest
Departement of Control Engineering & Information Technology
Pazmany P.S. 1/D
H-1111 Budapest
Hungary
Tel: +36 1 463-2030 or 463-3400
Fax: +36 1 463-2204
e-Mail: szirmay@iit.bme.hu
URL: www.fsz.bme.hu/˜szirmay/szirmay.html

Robert F. Tobler
VRVis Center for Virtial Reality and Visualization
Lothringer Str. 16
A-1030 Wien
Austria
Fax: +43-1-58801-18698
e-Mail: rft@cg.tuwien.ac.at
URL: www.cg.tuwien.ac.at

Dani Tost
Universitat Politecnica de Catalunya
Dept. L.S.I.

Planta 8, Despatx: 69
Ave. Diagonal 647 (Etseib)
E-08028 Barcelona
Spain
Fax: +34-93-401-60 50
e-Mail: dani@lsi.upc.es

Ingo Wald
Universität des Saarlandes
FR Informatik
Im Stadtwald
PF 15 11 50
D-66041 Saarbrücken
Germany
Tel: +49-681-302-3833
Fax: +49-681-302-3843
e-Mail: wald@graphics.cs.uni-sb.de
URL: graphics.cs.uni-sb.de/˜wald

Rüdiger Westermann
Universität Stuttgart
Institut für Informatik
Abt. Visualisierung und Interaktive Systeme
Breitwiesenstraße 20-22
D-70565 Stuttgart
Germany
Tel: +49-711-7816-368
Fax: +49-711-7816-340
e-Mail: ruediger.westermann@informatik.uni-stuttgart.de