

Report on Dagstuhl Seminar 9821
Hierarchical Methods in Computer Graphics
May 25 – 29, 1998

Organized by

Markus Gross, ETH Zürich, Switzerland
Heinrich Müller, University of Dortmund, Germany
Peter Schröder, Caltech Pasadena, USA
Hans-Peter Seidel, University of Erlangen, Germany

Over the last decade hierarchical methods, multiresolution representations and wavelets have become an exceedingly powerful and flexible tool for computations and data reduction within computer graphics. Their power lies in the fact that they only require a small number of coefficients to represent general functions and large data sets accurately. This allows compression and efficient computations. They offer both theoretical characterization of smoothness and coherence, insights into the structure of functions, and operators, and practical numerical tools which often lead to asymptotically faster computational algorithms. Examples of their use in computer graphics include

- curve, surface, and volume modeling,
- efficient triangle meshes, mesh simplification, subdivision surfaces,
- multiresolution surface viewing and automatic level of detail control,
- image and video editing, compression and querying,
- efficient solution of operators such as global illumination and PDEs as they occur in finite element modeling for animation and surgery simulation,
- flow and volume visualization.

There is strong evidence that hierarchical methods, multiresolution representations, and wavelets will become a core technique in computer graphics in the future.

This Dagstuhl Seminar has provided a forum for some of the leading researchers in this area to present their ideas and to bring together applications and basic research in order to exchange the requirements of systems, interfaces, and efficient algorithmic solutions to be developed. The seminar has been attended by 52 participants from 11 countries.

The main goal of the seminar has been to provide an opportunity for discussing ideas and work in progress. International conferences with their densely packed schedules usually leave little room for this sort of scientific exchange. Consequently, in order to save time for interaction and discussion, we have only scheduled 36 talks, and the unique atmosphere

of Dagstuhl has been immensely helpful to stimulate many inspiring discussions. The seminar has also benefited from the active participation of several young researchers. Here we are greatly thankful for the TMR (Training and Mobility of Young Researchers) funding provided by the European community. This funding has made it possible for several young researchers to attend the seminar and actively participate in the discussions.

The positive feedback that we have received after the seminar indicates that the workshop has been very well received, and we hope to be able to have a follow-up in the future.

Table of Contents

Paul Heckbert

Surface Simplification and Multiresolution Modeling

Leif P. Kobbelt

Multiresolution Modeling of Triangle Meshes

Reinhard Klein

Data Compression of Multiresolution Surfaces

Chandrajit Bajaj

Compression and Progressive Transmission of Meshes

Oliver Staadt

Multiresolution Meshing

Morten Daehlen

Hierarchical Structures for Terrain Visualization

Hans-Christian Hege

Enumeration of Symmetry Classes in Mesh Generation

Ken Joy

Robust Simplification of Tetrahedral Meshes

Bernd Hamann

Various Approaches to Hierarchical Data Modeling

Konrad Polthier

Interpolation of Triangle Hierarchies

Philipp Slusallek

Robust and Scalable Algorithms for Lighting Simulations

Marc Stamminger

Three-Point Clustering

Stephan Schäfer

Distributed Hierarchical Radiosity

Nelson Max

Hierarchical Rendering at LLNL

Alain Fournier

Real-time Rendering of Wavelet Compressed Light Fields

Alexander Keller

A Quasi-Monte Carlo Approach to Hierarchical Form Factor Calculation

Robert F. Tobler

Directional Importance for Hierarchical Stochastic Radiosity

Oliver Deussen

Multi-Resolution Modeling and Sketching of Plants

Larry L. Schumaker
On Local Bases for Bivariate Polynomial Spline Spaces

Tom Lyche
A Multiresolution Tensor Spline Method for Fitting Functions on the Sphere

Joerg Peters
Artesano: Hierarchical Splines for Modeling Surfaces

Franz-Erich Wolter
Cut Locus and Medial Axis in Global Shape Interrogation and Representation

Gregory Nielson
Cracking the Cracking Problem with Coons Patches

Joe Warren
Subdivision Schemes and Radial Basis Functions

Denis Zorin
Subdivision and Multiresolution Surface Representations

Nira Dyn
Computation of Normals to Surfaces Generated by Reversing Subdivision Rules

Heinrich Müller
Modeling with Subdivision Surfaces

Richard Bartels
Multiresolution Curves and Surfaces by Reversing Subdivision Rules

Bernd Girod
Hierarchical Image and Video Compression

Adi Levin
Combined Subdivision Schemes for the Design of Surfaces Satisfying Boundary Conditions

Michael Lounsbery
Subdivision Surfaces in Industry?

Dietmar Saupe
Optimal Hierarchical Space Subdivisions and Applications in Computer Graphics

Pere Brunet
Data Structures and Algorithms for Navigation in Highly Polygon-Populated Ship Environments

Werner Purgathofer
Levels of Detail for Animated Virtual Environments

Martin Roth
Volumetric Soft Tissue Modeling

Thomas Ertl
Hierarchical Methods in Volume Visualization

Surface Simplification and Multiresolution Modeling

Paul Heckbert, Michael Garland and Andrew Willmott
Carnegie Mellon University

It is easy to generate 3-D surface models consisting of millions of polygons, but such models are often too bulky for real time display, for storage, or for network transmission. Examples of voluminous data include terrains, outputs of laser rangefinders, and isosurfaces from volume visualization.

First I will present an algorithm for simplifying polygonal models by decimation. The algorithm repeatedly collapses edges of the model to reduce the number of polygons until the desired size or error is achieved. Errors are approximated in a fast, local manner using a quadric error metric. The algorithm yields high quality surface approximations more quickly than most previous algorithms.

I will also discuss an application of simplification techniques to the simulation of radiosity (the illumination of surfaces by other surfaces). Vertices or faces of a model are clustered to create a hierarchical, multiresolution model of a complex surface. This multiresolution model is then used directly in radiosity simulation, allowing faster solution than previous clustering techniques and reducing memory requirements significantly.

Multiresolution Modeling of Triangle Meshes

Leif P. Kobbelt
University of Erlangen

During the last years the concept of multi-resolution modeling has gained special attention in many fields of computer graphics and geometric modeling. In this paper we generalize powerful multi-resolution techniques to arbitrary triangle meshes without requiring subdivision connectivity. Our major observation is that the hierarchy of nested spaces which is the structural core element of most multi-resolution algorithms can be replaced by the sequence of intermediate meshes emerging from the application of incremental mesh decimation. Performing such schemes with local frame coding of the detail coefficients already provides effective and efficient algorithms to extract multi-resolution information from unstructured meshes. In combination with discrete fairing techniques, i.e., the constrained minimization of discrete energy functionals, we obtain very fast mesh smoothing algorithms which are able to reduce noise from a geometrically specified frequency band in a multi-resolution decomposition. Putting mesh hierarchies, local frame coding and multi-level smoothing together allows us to propose a flexible and intuitive paradigm for interactive detail-preserving mesh modification. We show examples generated by our mesh modeling tool implementation to demonstrate its functionality.

Data Compression of Multiresolution Surfaces

Reinhard Klein
University of Tübingen

In the talk we introduce a new compressed representation for multiresolution models (MRM) of triangulated surfaces of 3D-objects. Associated with the representation we present compression and decompression algorithms. Our representation allows to extract the surface at variable resolution in time linear in the output size. It applies to MRMs generated by different simplification algorithms like local vertex deletion or edge and triangle collapse. The time required to transmit models over communication lines and the space needed to store the MRMs is significantly reduced.

Compression and Progressive Transmission of Meshes

Chandrajit Bajaj
University of Texas, Austin

We present a topological ring layering scheme coupled with vector quantization for compressing both the topology (connectivity) and geometry (vertex coordinates) of arbitrary polygon meshes. The polygon mesh surface could be open or closed, non-manifold, and with multiple holes (arbitrary genus). The layered topological decomposition makes the compression and connectivity encoding as efficient as that for ribbon surfaces or triangle strip r . The vector quantization and geometry encoding of the vertex coordinates is done with novel error/distortion control parameters and allows progressive bit transmission as well as of the encoded connectivity information of the progressive simplified meshes. The separation of topology and geometry encoding permits all combinations of lossy or lossless topology and lossy or lossless geometry.

Multiresolution Meshing

Oliver Staadt
ETH Zürich

Methods to hierarchically represent triangular surface and tetrahedral volume data sets gain more and more importance. In this talk, I will present several vertex removal schemes, that have been developed in our group over the last few years. These methods are based on wavelet approximations using endpointinterpolating B-spline basis functions. The first

method is a quadtree-based triangulation where connectivity information is precomputed and stored in a look-up-table. The second method employs a data compression pipeline where both, triangular and tetrahedral meshes can be reconstructed progressively using Delaunay methods. It is further possible to extract high quality isolines and -surfaces. Finally, I will introduce extensions of progressive meshes to progressive tetrahedralizations and I will discuss problems such as tetrahedral folding or intersections.

Hierarchical Structures for Terrain Visualization

Morten Daehlen
SINTEF Applied Mathematics

With emphasis on digital terrain models, some major applications of hierarchical structures for topographic data were presented. Software issues and challenges with respect to handling massive data sets were also discussed. A combination of domain decomposition and hierarchical structures for terrain representation was used to speed up the visualization in a flight simulator for gliding aeroplanes.

Enumeration of Symmetry Classes in Mesh Generation

Hans-Christian Hege
Konrad-Zuse-Zentrum fuer Informationstechnik Berlin

Many algorithms in computer graphics process discrete configurations of elementary objects (pixels, vertices, cells, cell complexes, ...) which may assume discrete states of finite set (e.g. color, degree,...). Furthermore, often symmetries in space and state space are present. These induce equivalence classes, so-called orbits, in the set of all patterns. Usually algorithms are required to respect these symmetries, i.e. to treat patterns of the same orbit equally. For this and other reasons it is of interest to construct the set of all orbits for given patterns and symmetries. Even the knowledge of the cardinality of this set is useful, e.g. for considerations in algorithm design.

Examples are discussed and it is shown how the number of orbits can be computed in general. The procedure is exemplified by the patterns occurring in a generalized marching cubes algorithm for generating separating surfaces of a labeled voxel set. A formula derived by de Bruijn in algebraic combinatorics yields the number of orbits for different spatial symmetry classes and permutations in color space. Applying this to the example yields cardinalities which include as a special case the well-known numbers for the standard

marching cubes algorithm (for different symmetry groups). Finally constructive issues are addressed, like the generation of a complete set of representations and the search for the equivalence class which corresponds to some given pattern.

Robust Simplification of Tetrahedral Meshes

Ken Joy
University of California, Davis

We present a method for the construction of multiple levels of triangle and tetrahedral meshes approximating a bi- and trivariate functions at different levels of detail. Starting with an initial, high-resolution triangulation, we construct coarser representation levels by collapsing triangles in the two-dimensional case, and tetrahedra in three-dimensions. Each triangulation defines a linear spline function, where the function values associated with the vertices are the spline coefficients. Based on predicted errors, we identify the elements whose elimination would cause a minimal increase in error, and collapse them. Bounds are stored for individual elements and are updated as the mesh is simplified. We continue the simplification process until a certain error is reached. The result is a hierarchical data description suited for the efficient visualization of large data sets at varying levels of detail.

Various Approaches to Hierarchical Data Modeling

Bernd Hamann
University of California, Davis

Due to the ever increasing size of data in all science and engineering disciplines approaches are needed for their hierarchical representation (approximation) and visualization. One of the fundamental problems is the development of a "unifying," "universal" structure that allows to represent a hierarchy of massive data sets—both empirical and simulated ones—regardless of the original grid/mesh structure. We present various methods that seem promising in the context of developing such a general data format. We present approximation schemes allowing the hierarchical representation of univariate, bivariate, and trivariate data sets. In particular, we discuss possible solution avenues to the hierarchical approximation problem based on concepts known from best approximation, data-dependent triangulation, simulated annealing, tessellation (Voronoi diagrams), optimal knot (vertex) placement, and clustering.

Interpolation of Triangle Hierarchies

Konrad Polthier
TU Berlin

I consider the interpolation between two geometries, each represented as a hierarchical data structure, and impose a set of weak constraints to allow smooth interpolation. This approach works in the class of conforming triangulations.

Interpolation constraints: Let F be a family of triangle hierarchies.

1. The simplicial complex of the root triangles of each hierarchy in F is the same, i.e. for each pair of hierarchies $G, H \in F$ exists a bijective simplicial map Φ between the set of root triangles.
2. Each root triangle has a refinement edge, and Φ maps refinement edges to refinement edges.
3. Each hierarchy is refined using the Rivara Bisection algorithm, which refines triangles by bisecting the refinement edge.

Under these rather weak constraint there exist a smooth interpolation between hierarchies in F , and the hierarchy of the interpolation object is the union of the key hierarchies.

This concept has applications in animation of adaptively refined geometries, spline interpolation between animated objects, and in the interpolation in a more parameter family of hierarchies.

Robust and Scalable Algorithms for Lighting Simulations

Philipp Slusallek
University of Erlangen

Hierarchical methods have been introduced into the area of lighting simulation in the early nineties. They have reduced the quadratic time complexity of finite element style algorithms to a linear complexity through adaptive light transport on any level in the hierarchy and through clustering of surfaces. Nonetheless, in order to make these algorithms really interesting for commercial use, there remain several problems, in particular with respect to robustness and scalability.

In this presentation, two of these issues are addressed: Avoiding sampling problems while computing the form factor through the use of bounded computations and designing a new

iterative solution technique that avoids the need to store the links representing the linear system in a hierarchical radiosity setup.

Bounded computation avoid the need to use point-to-surface form factor sampling across a receiving patch in order to estimate the light transport. Instead interval arithmetic, bounding boxes, and cones-of-normals are used to compute conservative upper and lower bounds on the radiosity on the receiver (excluding visibility issues). A modified refinement oracle then takes advantage of this information and refines only in areas where the difference between the bounds is large. The result is a robust algorithm that can easily handle curved surfaces as well as clusters.

The Gauss-Seidel iteration scheme commonly used for solving the linear system in hierarchical radiosity setups requires the storage of all links, because all of them are reused in every iteration and recomputing them would be too expensive. By using a shooting algorithm modified for hierarchical scene descriptions there is no need to store all links. Instead, the available memory can be used to cache only important links that are likely to be reused later. A detailed error analysis ensure convergence, since the shooting scheme is not self-correcting as the Gauss-Seidel method. This approach makes it possible to compute radiosity solutions for very complex scenes while using only a fixed size cache for storing links.

Three-Point Clustering

Marc Stamminger
University of Erlangen

There has been great success in speeding up global illumination computation in diffuse environments. The concept of clustering allows radiosity computations even for scenes of high complexity. However, for lighting simulations in complex non-diffuse scenes, Monte-Carlo sampling methods are currently the first choice, because non-diffuse finite element approaches still exhibit enormous computation times and are thus only applicable to scenes of very modest complexity. In this talk we present a novel clustering approach for radiance computations, by which we overcome some of the problems of previous methods. The algorithm computes a radiance solution within a line space hierarchy, that allows us to efficiently represent light propagation and reflection between arbitrary non-diffuse surfaces and clusters.

Distributed Hierarchical Radiosity

Stephan Schäfer, Marco Zens, Dieter Fellner
University of Bonn

Today's most challenging application of computer graphics is the rendering of real scale scenes at photorealistic quality. As this task needs a lot of computing resources, numerous techniques have been applied to speed up rendering, one of these being parallel processing. This talk presents our approach to the efficient parallelization of Hierarchical Radiosity in a computer network. We introduced a slightly different formulation of the original sequential algorithm that minimizes the communication costs by exploiting the object oriented design of our rendering package. Additionally, we presented and discussed an adaptive load balancing, which incorporates individual CPU load.

Our approach to parallelize the rendering of a scene is to split the computation up into several jobs which are then distributed dynamically onto many clients by one distinguished computer, called server. This demand-driven client-server model offers simplicity, scalability and minimizes the need for communication, which, in the environment specified (LAN) is the bottleneck of every distributed computation.

The computation of the global illumination in a scene typically requires huge communication efforts. Our solution concentrates on the subtask that typically needs the most computing resources, i.e. the computation of the formfactors between patches, which is done by ray casting. Instead of launching rays through a permanently increasing amount of patches the visibility calculation is performed on the original scene objects only. This can be achieved by following the paradigm of keeping object information throughout the whole rendering process: each patch always knows its parent object, thus being able to call object specific class methods. This is the key to very small communication costs because there is no need to send the mesh to the clients.

To let all the clients finish their work nearly simultaneously an adaptive load balancing has been implemented. By observing the time actually needed by a client, permanently reducing jobsizes and allowing for distribution of jobs more than once if processors are idle nearly linear speedups for a moderate number of processors have been achieved.

Hierarchical Rendering at LLNL

Nelson Max, Mark Duchaineau, and Dan Schikore
Lawrence Livermore National Laboratory

This talk described several different projects. The first concerned image-based rendering of trees. Precomputed layered depth images of a hierarchy of twigs and branches, and

an RGBa texture of a single leaf, are adaptively selected according their distance from the viewpoint, and reprojected into the desired view. The per pixel data include color and surface normal, and are shaded after reprojection. One shading scheme is based on plane parallel radiance transport, which models the forest canopy as a volume density on infinitesimal surface scattering elements, distributed according to altitude, and polar angle between their surface normal and the zenith. Since this density of scatterers does not depend on x and y , the partial differential equations of radiance transport reduce to ordinary differential equations for radiance components in a collection of discretized radiance direction bins. Given the sky and sun illumination at the top of the canopy, and the BDRF of the ground as boundary conditions, these equations can be solved for the directional radiance distribution as a function of z , and used to shade the images.

Another project involved the visualization of a 3D variant of the Hilbert space-filling curve, using volume rendering. The variant is a closed curve, divided into segments with different colors and opacities, which move along the curve as the animation progresses. The volume rendering uses polyhedron projection in a front-to-back recursive traversal of an octree in the unit cube. If a cell has a constant color, it is projected and composited; otherwise it is divided into its eight subcubes for recursion.

An adaptive view dependent surface simplification for terrain maps was described, based on subdivision of isosceles right triangles in half by bisecting the right angle. If the subdivision levels of adjacent triangles differ by at most one, there will be no T joints. The subdivision proceeds based on screen projection error priority, until a triangle drawing budget is used up. For real time operation, the subdivision from the previous frame is incrementally modified using two queues, one of triangles to be subdivided, and one for triangle pairs to be rejoined.

Finally, a system for interactive viewing of a regularly sampled 3D scalar function, using color maps on the faces of an interactively specified cube was described. Sliders control the three moving slice planes and the time parameter, and the appropriate slices of the data are expanded from blocks of JPEG compressed section. Three orthogonal families of slice planes are required, giving an ultimate compression ratio of twenty to one, and real time software decompression of a 512 by 512 by 512 data set with 320 time steps was achieved using 16 processors.

Real-time Rendering of Wavelet Compressed Light Fields

Alain Fournier
University of British Columbia

In the past few years, we have acquired some experience in the acquisition, storage and manipulation of wavelet projections of multidimensional data. These are mostly light fields

(two space variables and two directional variables), bidirectional reflectance distribution functions (four directional variables), and radiance fields (two space variables and two directional variables) for light transport in Lucifer, our implementation of a light-driven global illumination algorithm. Most of this work is with Paul Lalonde and Bob Lewis.

I will present here a specific example, by Paul Lalonde and myself, where the goal is the real-time reconstruction of light-field objects compressed as wavelet projections.

Light field rendering techniques allow the rendering of objects in time complexity unrelated to their geometric complexity. The technique discretely samples the space of light rays exiting the boundary around an object and then reconstructs a requested view from these data. In order to generate high quality images a dense sampling of the space is required which leads to large data sets. These data sets exhibit a high degree of coherence and should be compressed in order to make their size manageable.

We present a wavelet-based method for storing light fields over planar domains. The parameterization is based on the Nusselt embedding, which leads to simplifications in shading computations when the light fields are used illumination sources. The wavelet transform exploits the coherence in the data to reduce the size of the data sets by factors of 100 or more without objectionable deterioration in the rendered images. The wavelet representation also allows a hierarchical representation in which detail can be added incrementally, and in which each coarser view is an appropriately filtered version of the finer detail.

Compression of wavelet coefficients is performed by thresholding the coefficients and storing them in a sparse hexadecary tree. The tree encoding allows random access over the compressed wavelet coefficients which is essential for extracting slices and point samples from the light field. The cost of reconstruction is linear in terms of the number of pixels displayed.

A demonstration will show real-time reconstruction of a 32x32x32x32 light-field object at about 8 frames/sec on a 166Mhz Pentium-based computer.

A Quasi-Monte Carlo Approach to Hierarchical Form Factor Calculation

Alexander Keller
University of Kaiserslautern

A new quasi-Monte Carlo approach for the calculation of hierarchical general form factors as used in hierarchical radiosity methods is introduced. The deterministic algorithm is based on ray shooting and low discrepancy sampling, runs in subquadratic time concerning the number of elements in the scene, and is superior to previous Monte Carlo algorithms. The algorithm also reveals some disadvantages of kernel discretization. Comparing the work

to be done by a quasi-random walk for a hierarchical solution discretization of the same quality, the question raises, whether it is necessary to discretize the kernel at all.

Directional Importance for Hierarchical Stochastic Radiosity

Robert F. Tobler and Werner Purgathofer
Vienna University of Technology

Stochastic radiosity normally operates on a premeshed scene. Since this is not desirable, as the places with high variation in the radiosity function are not known in advance, we introduce a hierarchical extension to standard radiosity methods, that subdivides the input surfaces as the simulation progresses. The main idea is to track the radiosity at different levels of resolution, and subdivide in places where these representations differ significantly. We also show how this scheme can be extended to CSG models, by either operating in the (u,v) -space of the primitives, or by building a three-dimensional storage scheme for radiosity.

There is however an efficiency problem: in order to subdivide in locations of interest, a lot of photons have to be simulated that arrive at these locations. Other regions of little interest will thereby also receive a huge number of photons. Thus it is desirable to build a four-dimensional data structure (2 spatial and 2 directional dimensions) in order to store directional importance and direct photons to places of interest. We are currently investigating different schemes for storing this information adaptively.

Multi-Resolution Modeling and Sketching of Plants

Oliver Deussen
University of Magdeburg

A modeling method for creating natural branching structures such as plants is presented. Structural and geometrical information is encapsulated in objects that are combined to a graph that forms the description of the model. Global and partial constraint techniques are integrated on the basis of tropisms, free-form deformations and pruning operations allow modeling of specific shapes. The models are used to generate whole ecosystems by combining several species. We show how these systems can be generated graphically or by simulation. Geometry is reduced by approximative instancing, the scenes are rendered either by parallel raytracing or by a hardware-supported ray casting algorithm.

On Local Bases for Bivariate Polynomial Spline Spaces

Larry L. Schumaker and Oleg Davydov
Vanderbilt University

Given a regular triangulation Δ , we consider the space

$$\mathcal{S}_d^r(\Delta) := \{s \in C^r(\Omega) : s|_T \in \mathcal{P}_d \text{ for all triangles } T \in \Delta\},$$

where \mathcal{P}_d is the space of polynomials of degree d , and Ω is the union of the triangles in Δ . Bivariate polynomial spline spaces have been heavily studied in recent years. Among other things, there is an extensive theory of dimensions, bases, and approximation power. In addition to fixed triangulations, we are also interested in *scales of spline spaces* defined on a sequence of triangulations. Such sequences of spline spaces arise through the process of *refinement* of the triangulation, leading to a nested sequence of spline spaces, and are important in a variety of multiresolution and hierarchical applications.

In the univariate case it is easy to construct spline bases which have a variety of nice properties such as local support, stable representation of polynomials, and local linear independence. While locally supported bases have been constructed for bivariate splines for the case $d \geq 3r + 2$, the existing constructions in the literature do not yield stable local representations of polynomials. In this paper we show how to modify existing construction methods (based on a mix of cofactor and Bernstein-Bézier techniques) in order to get stable bases. We also show how to construct bases which are locally linearly independent (which is important for certain almost interpolation problems).

A Multiresolution Tensor Spline Method for Fitting Functions on the Sphere

Tom Lyche, University of Oslo
Larry L. Schumaker, Vanderbilt University

We show how to use multi-resolution methods with tensor products of polynomial splines and trigonometric splines to fit data on the sphere. The method produces surfaces which are tangent-plane continuous, and provides a convenient data compression algorithm for dealing with large amounts of data.

Artesano: Hierarchical Splines for Modeling Surfaces

Joerg Peters and C. Gonzales
Purdue University

Artesano is a hierarchical environment for conceptual 3D modeling. It is based on surface splines, a representation for C^1 manifolds that generalizes B-splines to meshes with polyhedral connectivity, allowing n -valent mesh points and m -sided facets. Features include direct surface manipulation and a lean interface allowing for subdivision of the control structure for localized change, extrusions and change of genus.

Cut Locus and Medial Axis in Global Shape Interrogation and Representation

Franz-Erich Wolter
University of Hannover

The cut locus C_A of a closed set A in the Euclidean space E is defined as the closure of the set containing all points p that have at least two shortest segments to A . We present a theorem stating that the complement of the cut locus i.e. $E \setminus (C_A \cup A)$ is the maximal open set in $(E \setminus A)$ where the distance function with respect to the set A is continuously differentiable. This theorem includes also the result that this distance function has a locally Lipschitz continuous gradient on $(E \setminus A)$.

The medial axis of a solid D in E is a subset of D containing all points being center of a disc of maximal size that fits in the domain D . We associate with the medial axis of a domain D the maximal disc radius function assigning to a medial axis point p the radius of the maximal disc with center p . We assume in the medial axis case that D is closed and that the boundary ∂D of D is a topological (not necessarily connected) hypersurface of E . Under these assumptions the medial axis of D equals that part of the cut locus of ∂D which is contained in D . The main result states that the medial axis has the homotopy type of its reference solid if the solid's boundary surface fulfills certain regularity requirements. The medial axis with its related maximal disc radius function can be used to reconstruct its reference solid D because D is the union all maximal discs that fit in D . Keeping the medial axis of a reference solid D fixed and modifying the associated disc radius function e.g. by shrinking or expanding the maximal disc radius function for some subsets of the medial axis yields a natural design tool allowing in a simple way global shape modifications like slimming or fattening the shape.

The cut locus concept offers a common frame lucidly unifying different concepts such as Voronoi diagrams, medial axes and equidistantial point sets. In this context we explain

that the equidistantial set of two disjoint point sets is a subset of the cut locus of the union of those two sets and that the Voronoi diagram of a discrete point set equals the cut locus of that point set. We present results which imply that a non-degenerate C^1 -smooth rational B-spline surface patch which is free of self-intersections avoids its cut locus. This implies that for small enough offset distances such a spline patch has regular smooth offset surfaces that are diffeomorphic to the unit sphere. Any of those offset surfaces bounds a solid (which is homeomorphic to the unit ball) and this solid's medial axis is equal to the progenitor spline surface. The spline patch can be manufactured with a ball cutter whose center moves on the regular offset surface and the radius of the ball cutter equals the offset distance.

Cracking the Cracking Problem with Coons Patches

Gregory Nielson
Arizona State University

Volume modeling involves the determination of a mathematical model for volume data. Many new sensors and simulation techniques are now producing volume data which consists of locations (x, y, z) and associated dependent data values ρ . The list of tuples (x_i, y_i, z_i, ρ_i) is modeled with a trivariate function $F(x, y, z)$. The model, F , has the potential to be a compact representation of the volume data allowing for analysis and visualization. The spatial distribution of the sample locations affects the form and methods of determining the modeling function F . In this research project we are investigating the merits of adaptive, least squares fits to noisy, redundant data associated with 3D ultrasound sensors. Adaptive techniques in dimensions higher than one exhibit the so called "cracking problem". In this presentation we will survey some of the previous attempts of solving the cracking problem and a novel approach based upon the transfinite interpolant ideas of the "grandfather" of computer graphics, Steven A. Coons.

Subdivision Schemes and Radial Basis Functions

Joe Warren
Rice University

The speaker discusses the relationship between subdivision schemes, variational problems and radial basis. Given a differentiation operator Δ and a radial basis function $\Psi(t)$ the speaker describes a change of basis $\Phi(t) = \sum d(i)\Psi(t - i)$ where $d(i)$ are a discrete approximation to Δ . If $\tilde{d}(x)$ is the generating function of the form $\tilde{d}(x) = \sum d(i)x^i$, then $\Phi(t)$ has a subdivision scheme with generating function $\tilde{d}(x^2)/\tilde{d}(x)$.

Subdivision and Multiresolution Surface Representations

Denis Zorin
Stanford University

We discuss how multiresolutional representations of surfaces can be constructed using subdivision. Our representation is a pyramid representation suitable for independent manipulation of geometry at different levels of resolutions. The hierarchical structure allows one to use highly efficient adaptive algorithms for modification and rendering of the models.

In our representation subdivision plays a role similar to the role of the "lowerpass" synthesis filter in wavelet representations. While properties of filters and related bases (approximation, regularity, stability etc.) are well understood in the functional setting on regular grids, only C^1 -continuity was explored in sufficient detail for surface subdivision. While similar properties can be defined and studied for bases and frames used for representing surfaces, due to the fundamental difference between functions and surfaces, even finding suitable definitions of such properties as approximation remains an open question.

Computation of Normals to Surfaces Generated by Reversing Subdivision Rules

Nira Dyn, D. Levin, and P. Shenkman
Tel Aviv University

Explicit formulae for normals of surfaces generated by the Butterfly subdivision scheme at the control points of each subdivision level are presented. These formulae compute the normals at a control point in terms of that control point and a finite number of neighboring control points. There is one formula for all vertices of any valency between 4 to 10. The formula for valency 6 is derived explicitly since most vertices after one subdivision iteration have this valency. For the other valencies the formula involves a root of a cubic equation with coefficients depending on the valency. The method for deriving this formula involves the eigenvectors of a matrix which describes the subdivision scheme locally near a control point. This method can be used for deriving the formulae for normals of C^1 -surfaces generated by any subdivision scheme.

The computed normals are used for Gouraud shading, which enables high quality rendering after a small number of subdivision iterations (two instead of five). Another application of the normals is to the computation of an offset surface of a Butterfly scheme surface and its rendering with Gouraud shading.

Modeling with Subdivision Surfaces

Heinrich Müller, Markus Kohler, Reinhard Jaeschke
University of Dortmund

Most of the work on subdivision curves and surfaces concerns mathematical aspects like convergence or degree of smoothness. In contrast, we focus on computational aspects like space and time requirements of the calculation of approximating polygonal chains or meshes. Subdivision curves and surfaces result as limit shape of an infinite sequence of polygonal chains or meshes, respectively, by iteratively calculating finer meshes according to particular rules. Well-known schemes on which our investigations are based are those of Chaikin, DLM (Dyn et al.), Doo/Sabin, and Catmull/Clark.

The first result is an alternative scheme of approximate mesh calculation. It differs from the usual breadth-first strategy of evaluating the meshes level by level by a new depth-first approach which refines the meshes simultaneously on all levels and only maintaining a path of the dependency dag. In this way, the working memory is reduced to a practical amount even for higher levels of iteration.

The second result is an extension of known subdivision schemes which is composed of several levels of refinement, hence yielding an adaptive level-of-detail representation. It turns out that the modified schemes have the property that any configuration can be calculated from any other one. The basic observation is that the refinement operators are invertible, also in the case of partial refinement. This could be shown for Chaikin and DLM subdivision curves. For meshes, this property holds at least for faces and stars, respectively, up to a number of points interesting for practical applications. For an arbitrary number it remains an open problem.

Multiresolution Curves and Surfaces by Reversing Subdivision Rules

Richard Bartels
University of Waterloo

Subdivision rules are applicable to surfaces represented as a mesh of points connected by edges. A subdivision rule for a curve or surface is given as a formula for replacing given points and edges by newer, more numerous ones. Each subdivision rule is applicable to only certain topologies of mesh.

This talk concentrates on rules for curves and tensor product surfaces, where the rules are most typically given in the form of a matrix that transforms "course points" into "fine points". We use the matrix to make some observations on the connection between

subdivision rules, the underlying scale and wavelet functions implied by the rules, and the inner product used to define orthogonality. We show how the inner product influences the support of the wavelets. We argue that, for surfaces in graphics, the conventional inner product is not as suitable as one based upon discrete least squares, and we show a construction that uses a local least squares inner product to provide finite analysis and reconstruction processes for subdivision surfaces. All examples involve Chaiken's rule but have been carried out for a number of other subdivisions.

Hierarchical Image and Video Compression

Bernd Girod
University of Erlangen

Hierarchical image and video coding schemes offer both excellent coding efficiency and the ability to support scalability. This talk gives an introduction to the principles of multiresolution coding of image and video signals. For Gaussian stationary random processes, rate distortion theory suggests that optimum compression can be achieved by independent encoding of frequency components. The characteristic shape of the power spectrum of images makes resolution pyramids particularly suitable. Besides critically sampled subband pyramids, that include the Discrete Wavelet Transform, oversampled pyramid decompositions are discussed. Oversampled subband pyramids are often better suited for scalable coding schemes.

Combined Subdivision Schemes for the Design of Surfaces Satisfying Boundary Conditions

Adi Levin
Tel Aviv University

We present an extended notion of subdivision schemes, where boundary conditions on the limit surfaces are considered at every step of the subdivision. The resulting limit surfaces satisfy exactly the given boundary conditions, no matter how they are represented.

This results in simple algorithms for creation of surfaces satisfying boundary conditions.

Our analysis is not restricted to a specific scheme nor to a specific kind of boundary condition, and we prove the conditions for smoothness of the limit surfaces.

Subdivision Surfaces in Industry?

Michael Lounsbery
Alias|Wavefront, Seattle

In this talk, we examine whether subdivision surfaces are appropriate for industrial use. Design and entertainment are two large branches of industry, and each has very different characteristics that can influence how helpful subdivision surfaces may be in it, or how likely they are to be accepted. We examine obstacles to their acceptance in the different branches, and discuss issues of which version of subdivision surfaces may be most useful.

The areas of entertainment and design differ greatly. Design is a process of communication. It begins with conceptual design, where artistic designers sketch out many possibilities for the design. The next stage is modeling, where more technical people precisely model the designers' shape in the computer. The exacting engineering demands of functionality and manufacturability have a subsequent impact on the results, often requiring redesign.

In entertainment, a precise 3D model is far less important – what is wanted is a good look in a rendered image. The entertainment process tends to be much faster than the design process, and people in the entertainment industry are far more open to innovations in process than are those in design.

Obstacles to subdivision surfaces in design include that they don't yet fit well with existing standards, and that a comprehensive package is needed before they can realistically be accepted. There is also a perception of "weirdness" associated with subdivision surfaces, where they are often seen as academic and non-practical. The real relevance and practicality of subdivision surfaces for solving important issues of continuity and topology needs to be stressed.

Optimal Hierarchical Space Subdivisions and Applications in Computer Graphics

Dietmar Saupe
University of Leipzig

Space subdivisions can be applied to generate region-based approximation of functions or images and for speed-up of search processes such as finding objects that intersect a given ray. With each space partitioning there is an associated cost (e.g. storage, space) and another functional such as approximation error or expected computation time. The optimization problem consists of finding the partitioning which yields the best quality for a given cost budget. When the partitioning is hierarchical and some technical conditions on the cost and quality functionals are given the optimization problem can be solved by the generalized BFOS algorithm of Chon, Lookabaugh, Gray. We review this algorithm and present two applications:

1. Memory constrained isosurface cell extraction.

As a result we obtain a method which optimally trades memory for speed and is a generalization of the octree method of Wilhelms, Van Gelder.

2. Region-based (fractal) image compression.

In this application the state-of-the-art greedy HV-decomposition is replaced by corresponding optimal subdivisions at some additional preprocessing cost.

Data Structures and Algorithms for Navigation in Highly Polygon-Populated Ship Environments

Pere Brunet, Carlos Saona, and Isabel Navazo
Universidad Politecnica de Cataluna

A pre-processing visibility algorithm for navigation in very complex virtual ship environments is presented. The algorithm computes a weak invisibility set for every node in an octree decomposition of the scene. Octree subdivision is performed in zones of non uniform visibility. Node visibility is computed in terms of point visibilities, using a set of potential occluder objects. The occluder set is automatically increased when too small node invisible sets are detected. Octree coherence is used in order to avoid duplication of computations. The weak visibility graph is a directed graph that connects every node of the visibility octree with the set of invisible nodes in a hierarchical space decomposition of the scene objects. This weak visibility graph is used in conjunction with a multiresolution object representation. The algorithm performance is discussed through different simulations.

Levels of Detail for Animated Virtual Environments

Werner Purgathofer, Dieter Schmalstieg
Vienna University of Technology

To achieve a constant frame rate in real-time rendering, image fidelity is traded for speed by modeling objects at multiple Levels Of Detail (LOD). Instead of pre-modeling a few discrete LODs, recent approaches use Smooth Levels of Detail (SLOD) - which allow an almost continuous representation of geometry. SLODS can also be used in distributed virtual environments for progressive transmission of geometry. Large unstructured geometry models can also be preprocessed for viewpoint-dependent adaptive selection of geometric detail.

Levels of detail for animated virtual environments improve SLOD methods in three areas: (1) A geometry model can be decomposed into smaller regions, for which a coarse view-dependent detail selection can be made with considerably reduced computational effort. (2) This decomposition is compatible with hierarchical scene graphs commonly used in virtual environment modeling. (3) The hierarchical model is combined with a method for real-time skeleton animation. This approach allows large scenes composed of deformable objects to be used with real-time rendering.

Volumetric Soft Tissue Modeling

Martin Roth, ETH Zürich

In this talk a Finite Element approach for volumetric soft tissue modeling in the context of facial surgery simulation is presented. We elaborate on the underlying physics and address some computational aspects of the finite element discretization.

In contrast to existing approaches speed is not our first concern, but we strive for the highest possible accuracy of simulation. We therefore propose an extension of linear elasticity towards incompressibility and nonlinear material behavior in order to describe the complex properties of human soft tissue more accurately. Furthermore, we incorporate higher order interpolation functions using a tetrahedral Bernstein-Bezier formulation, which has various advantageous properties.

Experimental results obtained from a synthetic block of soft tissue and from the Visible Human Data Set illustrate the performance of the envisioned model.

Hierarchical Methods in Volume Visualization

**Thomas Ertl
University of Erlangen**

Volume rendering and isosurface extraction from large 3D cartesian datasets are two visualization methods where hierarchical approaches can be applied successfully. First we presented sparse grids as a method for representing a function in a highly compressed manner with only moderately increased interpolation error. For the actual volume rendering of the sparse grids we use the combination technique, which allows us to employ hardware assisted trilinear interpolation of the 3D texture mechanism of OpenGL since it works on a linear combination of uniform grids. The second part of the talk focussed on adaptive isosurface extraction from hierarchically refined unstructured grids. Based on a fast reconstruction algorithm we can refine the grid from the compressed representation to an arbitrary level. Isosurfaces extracted from one level of an affine hierarchy can be refined to the next level by a fast algorithm which allows for progressive transmission to a Java applet which incrementally updates the VRML scene graph.