

Conference/6 11 August 1995 Exhibition/8 10 August 1995 Los Angeles Convention Center Los Angeles California USA

S I G G R A P H 9 5

Interactive Walkthrough
of Large Geometric
Databases

C o u r s e

32

N O T E S

22nd International Conference on Computer Graphics and Interactive Techniques

Interactive Walkthrough of Large Geometric Databases

Organizer
ERIC BRECHNER
Boeing Computer Services

Lecturers
THOMAS FUNKHOUSE
AT&T Bell Laboratories

NED GREENE
Apple Computer Inc.

JAMES HELMAN
Silicon Graphics Inc.

JAREK ROSSIGNAC
IBM TJ Watson Research Center

C o u r s e

32

NOTES



SIGGRAPH 1995

22nd International Conference on Computer
Graphics and Interactive Techniques

Conference/6-11 August 1995
Exhibition/8-10 August 1995
Los Angeles Convention Center
Los Angeles California USA

Interactive Walkthrough of Large Geometric Databases

Course Notes for SIGGRAPH '95

Course Organizer

**Eric L. Brechner
The Boeing Company**

Course Speakers

**James Helman
Silicon Graphics Inc**

**Ned Greene
Apple Computer**

**Jarek Rossignac
IBM T J Watson Research Center**

**Thomas A. Funkhouser
AT&T Bell Laboratories**

Abstract

This course will focus on techniques, algorithms, data-structures and databases for displaying very large geometric databases interactively (greater than 1 million polygons drawn at least 10 frames per second). Beginning with a discussion of basic techniques and leading to state-of-the-art algorithms, the speakers will address key issues in walkthrough including visibility computations, automatic object simplification, and memory management through database subset pre-fetching. Speakers will show real applications of these algorithms to a variety of areas including visual simulation, virtual reality, architecture, and digital mockup.

Attendees of this course will better understand the key issues in dealing with very large geometric databases. They will be provided basic and state-of-the-art techniques to overcome hardware and software limitations that preclude the high frames rates necessary for interactive inspection of complex geometric scenes. These techniques will be illustrated by real examples of working walkthrough applications using databases with between 1 million and 500 million polygons.

A working knowledge of interactive computer graphics, including the mechanisms of matrix transformation, perspective, and raster graphics, will form a good foundation for the course material.

Course Schedule

A Introduction - 10 minutes

Statement of the problem Introduction of the speakers Agenda. - Eric Brechner

B Graphics Techniques for Walkthrough Applications - 75 minutes

An introduction to techniques common to most walkthrough applications with examples from the IRIS Performer toolkit. The methods include multi tasking
{ view frustum culling level-of-detail frame rate control, and database paging }
{ James Helman }

- Morning Break

C Hierarchical Visibility and Tiling 75 minutes

Hierarchical approaches to accelerating visibility computations in extremely complex scenes { Ned Greene }

- Lunch

D Geometric Simplification - 75 minutes

The generation and exploitation of multi-resolution graphic models for the interactive visualization of complex mechanical or architectural 3D scenes { Jarek Rossignac }

- Afternoon Break

E Database Management - 75 minutes

Algorithms for computing and pre-fetching a small subset of a disk-resident database to store in memory during an interactive walkthrough { Thomas Funkhouser }

- Seventh inning Stretch

F Wrap-up and Future Directions - 20 minutes

Hierarchical level of detail culling and data staging for interactive walkthrough of infinite databases, including a 3D Web browser { Eric Brechner }

2:45

3:00 ~~2:45~~

4:40

5:00

Contents

Abstract	iii
Course Schedule	v
About the speakers	xvii
A Introduction	A-1
1 Slides	A-3
2 Introduction	A-7
2.1 Statement of the problem	A 7
2.2 Overview of the notes	A 8
B Graphics Techniques	B-1
3 Graphics Techniques	B-3
3.1 Introduction	B-3
3.2 Motivation of Performance Requirements	B-4
3.2.1 Visual Acuity	B-4
3.2.2 Field of View	B-5
3.2.3 Latency	B-5
3.2.4 High Frame Rate	B-5
3.2.5 Constant Frame Rate	B 7
3.3 Performance in Graphics Pipelines	B-7
3.4 Reducing Geometric Complexity	B-8
3.4.1 Texture Mapping	B-8
3.4.2 Billboards	B-9
3.4.3 Level of Detail	B-10
3.4.4 Geometric	B 10
3.4.5 Texturing	B 10
3.5 Optimizing Runtime Rendering	B 10
3.5.1 Visibility Culling	B-10

3 5 2	Level of Detail Switching	B-12
3 5 3	Maintaining Constant Frame Rate	B 12
3 5 4	Parallel Processing	B 13
3 5 5	Mode Sorting	B 14
3 5 6	Latency Control	B 14
3 6	Optimizing the Graphical Database	B 16
3 6 1	Triangle Meshes	B 16
3 6 2	Spatial Organization	B 16
3 6 3	Graphics Attributes	B 17
3 6 4	Eliminating Transformations	B 17
3 6 5	Precomputed Animation	B 17
3 7	Database Examples	B 17
3 7 1	Architectural Exteriors	B 17
3 7 2	Architectural Interiors	B 17
3 8	Handling Very Large Databases	B 18
3 8 1	Paging Geometry and Texture from Disk	B 18
3 8 2	Paging Texture to Graphics Subsystem	B 20
3 9	Beyond the Rendering	B 20
3 9 1	Interface with M/CAD Systems	B 20
3 9 2	Collision Detection	B 21
3 9 3	Networking	B 21
3 10	Conclusions	B 21
C	Hierarchical Visibility and Tiling	C-1
4	Introduction	C 3
5	Hierarchical Z-Buffer Visibility	C 5
5 1	Overview	C 5
5 2	Introduction	C 5
5 3	The Hierarchical Z Buffer Visibility Algorithm	C 7
5 3 1	The Object-Space Octree	C 7
5 3 2	The Image Space Z-Pyramid	C 11
5 3 3	The Temporal Coherence Scheme	C 15
5 4	Building and Maintaining The Octree	C 15
5 4 1	Building an Octree from Unorganized Primitives	C 16
5 5	Implementation and Results	C 18
5 5 1	Parallel Performance	C 23
5 5 2	Use of Graphics Hardware	C 24
5 6	Conclusion	C 25

6 Hierarchical Polygon Tiling	C-27
6 1 Overview	C-27
6 2 Introduction	C-28
6 3 Previous Work	C-30
6 3 1 Warnock Subdivision	C-30
6 3 2 Coverage Masks	C-30
6 4 Triage Coverage Masks	C-31
6 4 1 Triage Edge Masks	C-32
6 4 2 Compositing Triage Masks	C-32
6 4 3 Building Lookup Tables	C-35
6 5 The Hierarchical Tiling Algorithm	C-36
6 5 1 Data Structures	C-36
6 5 2 The Basic Hierarchical Tiling Algorithm	C-37
6 6 Hierarchical Object-Space Culling	C-42
6 7 Implementation and Results	C-43
6 8 Conclusion	C-45
D Geometric Simplification	D-1
7 Geometric Simplification	D 3
7 1 Introduction	D-3
7 2 Rendering cost	D-4
7 3 Graphics acceleration techniques	D 5
7 3 1 Meshing or storing transformed vertices	D-5
7 3 2 Frustum culling	D-6
7 3 3 Smart caching and pre-fetching	D-6
7 3 4 Visibility	D-6
7 3 5 Polygon count reducing techniques	D-7
7 4 Geometric simplification	D-8
7 4 1 Grading	D 8
7 4 2 Triangulation	D 8
7 4 3 Clustering	D-8
7 4 4 Synthesis	D-9
7 4 5 Elimination	D 9
7 4 6 Adjustment of normals	D-9
7 4 7 Exploitation	D-9
7 4 8 Advantages	D-10
7 5 Experiments	D-10
7 6 Conclusion	D-11
E Database Management	E-1
8 Slides	E-3

9 Database Management	E 15
9 1 Introduction	E 16
9 1 1 Previous Work	E 17
9 1 2 Goals	E 19
9 1 3 Problem Statement	E 23
9 1 4 System Organization	E 24
9 2 Modeling	E-24
9 2 1 Model Loading	E-24
9 2 2 Model Representation	E-26
9 2 3 Object Abstraction	E-27
9 2 4 Results	E 28
9 3 Precomputation	E-29
9 3 1 Spatial Subdivision	E-29
9 3 2 Visibility Precomputation	E 30
9 3 3 Results	E 34
9 4 Interactive Walkthrough	E 34
9 4 1 User Interface	E 36
9 4 2 Display Management	E 37
9 4 3 Memory Management	E-43
9 5 Results	E-48
9 5 1 Display Management	E-48
9 5 2 Memory Management	E 52
9 6 Conclusion	E-54
9 7 Acknowledgements	E 55
F Wrap-up and Future Directions	F-1
10 Slides	F 3
11 Wrap-up and Future Directions	F 7
11 1 The infinite database problem	F 7
11 1 1 Four key concepts	F 7
11 1 2 Choosing the right approximation	F 8
11 1 3 A possible approximation choice	F 9
11 2 The ultimate 3D Web browser	F 10

List of Figures

3 1	Sources of latency in a hypothetical virtual environment system	B-6
3 2	Multiple image artifact when fixating on moving object.	B-7
3 3	View into a simple town using texture mapping and billboards	B-9
3 4	Transition from a higher level-of-detail model to a lower one using fade (top) and morphing (bottom)	B-11
3 5	Graphics Load Management.	B-13
3 6	Pipelined processing model for multiprocessor graphics workstations	B-15
3 7	Trading off throughput for latency in pipeline multiprocessor	B-16
3 8	Visualization of a building used for city zoning discussions	B-18
3 9	Building interior using radiosity solution no texture mapping	B 19
3 10	Architectural scene using radiosity solution and texture mapping	B-19
5 1	If a cube is hidden then all geometry it contains is also hidden	C-8
5 2	Ordering of octants by visibility priority	C-10
5 3	A primitive inside a visible cube is 'nearly visible'	C-11
5 4	A scene and its corresponding z pyramid	C-13
5 5	Does a z-pyramid hide a primitive?	C-14
5 6	An office environment rendered with hierarchical visibility	C-20
5 7	Top view of model space showing viewing frustum and octree subdivision	C-21
5 8	Terrain models rendered with hierarchical visibility	C-22
5 9	Parallel performance graph	C-23
6 1	Triage coverage masks classify subcells as inside outside, or intersecting an edge	C-31
6 2	Constructing a triage coverage mask for a convex polygon from the triage masks of its edges	C-33
6 3	Compositing triage coverage masks	C-34
6 4	With triage coverage masks classification of subcells as <i>covered</i> or <i>vacant</i> is definitive	C-36
6 5	Schematic diagram of a mask pyramid	C-38
6 6	PSEUDOCODE LISTING	C-40
6 7	Antialiased frame rendered with the hierarchical polygon tiling algorithm	C-44
9 1	Exterior view of Soda Hall	E-20

9 2	Exterior view of Soda Hall cut open by a horizontal plane at the sixth floor	E 20
9 3	Typical office with furniture	E 20
9 4	Board room with furniture	E 20
9 5	Radiosity rendering of hallway	E-22
9 6	Radiosity rendering of office	E-22
9 7	Radiosity computation and results are independent of observer viewpoint	E 22
9 8	Polygonal mesh generated for the office during radiosity computation	E 22
9 9	System overview	E 24
9 10	Loading operations of the modeling phase	E 25
9 11	Three LODs for a chair	E-26
9 12	Object instances can share geometries stored in an object definition	E 27
9 13	An object instance can store its own geometries	E 27
9 14	Different door handle representations constructed using a parameterized generator	E 27
9 15	Three LODs for a canary	E 28
9 16	Functional steps of the precomputation phase	E 30
9 17	Spatial subdivision of the sixth floor of Soda Hall The image on the left shows the actual three dimensional subdivision while the image on the right shows a two dimensional schematic representation	E 31
9 18	A sight line stabbing a portal sequence	E 31
9 19	Bowtie shaped region containing sightlines stabbing a portal sequence (supple gray) Opaque boundaries are shown in solid black Extremal sightlines are shown as dashed lines	E 32
9 20	Cell visibility (supple gray) for source cell (dark gray)	E 32
9 21	Two dimensional schematic diagram of a) cell to-cell visibility (supple gray) and b) cell to-object visibility (shown as solid black squares) for a particular source cell (dark gray)	E-33
9 22	Cell visibility in three dimensions Visible region (beams) and cell to-cell visibility (outlined) are shown for one source cell	E 33
9 23	Organization of data in the display database	E 35
9 24	Functional operations of the walkthrough phase	E 36
9 25	View frustum variables	E-37
9 26	Mapping view frustum to perspective projection	E 37
9 27	Visualization program supports two views	E-37
9 28	Panels used for controlling parameters for a) observer navigation b) visibility determination and c) memory management	E-38
9 29	Visible region is a wedge that typically narrows as it traverses through more portals	E-40
9 30	Visible region for view frustum in two dimensions	E-40
9 31	Objects in the eye-to-object visibility are shown as by solid squares	E-40
9 32	All objects incident upon cells in the eye to-cell visibility	E-41
9 33	Objects in the eye-to-object visibility	E-41
9 34	The observer range cells (cross hatch) contain all observer view positions possible during the upcoming $N = 4$ frames Each cell is labeled by the number of frames before the observer can be resident in it	E-44

- 9 35 The lookahead cells (stipple gray) contain all objects that can be visible to the observer during the upcoming N frames. Each cell is labeled by the number of frames before it can become visible to the observer. E-45
- 9 36 Lookahead objects are stored in memory only up to the LOD at which they can possibly be rendered during the next N frames. Each cell is labeled and shaded by the maximum level of detail any object incident upon it is stored in memory – darker shades of gray represent higher levels of detail. E-46
- 9 37 Cache management algorithm results. Each cell is labeled by the number of frames since objects incident upon it were included in the lookahead set. Shading for each cell indicates whether or not it contains objects in the memory resident cache (stipple gray) read set (left hatch) and resident set (right-hatch). E-47
- 9 38 Test observer path through the top four floors of Soda Hall. E-49
- 9 39 Frame time for each observer viewpoint along the entire test observer path during the test using both visibility determination and detail elision. E-50
- 9 40 Images of library generated using (left) no detail elision (19 821 polygons) and (right) using the *Optimization* detail elision algorithm with a target frame time of 0.05 seconds (3 568 polygons). E-51
- 9 41 Rendering times during *Detail Elision* test for each observer viewpoint along the test observer path. E-52
- 9 42 Compute time and rendering time during *Detail Elision* test for each observer viewpoint along the portion of test observer path where computation was most expensive. E-52
- 9 43 Frame time for each observer viewpoint along the entire test observer path during the test with memory management. E-53
- 11 1 Approximating a cluster of polygons. F-10
- 11 2 A possible location scheme based on Internet address. F-11

List of Tables

6 1	Rules for compositing subcells of triage masks	C-33
9 1	Multi resolution modeling statistics for Soda Hall	E-29
9 2	Mean and maximum statistics for cells in the spatial subdivision of Soda Hall	E-34
9 3	Mean and maximum compute time rendering time and frame time statistics collected during display management tests	E-50
9 4	Mean and maximum numbers of cells objects and polygons rendered during display management tests	E-50
9 5	Mean and maximum compute time rendering time swap time, and frame time statistics collected during tests with and without memory management.	E-53
9 6	Mean and maximum set statistics collected during memory management tests	E-53

About the speakers

Eric L. Brechner

Senior Principal Scientist
The Boeing Company
PO Box 24346 MS 7L-43 Seattle WA 98124-0346
*brechner@boeing.com*¹

Eric Brechner is a Senior Principal Scientist in the Research and Technology Division of Boeing Computer Services. He has been with Boeing for five years working in the areas of large scale visualization, computational geometry, network communications, data-flow languages, and software integration. He was the principal architect of FlyThru™, the walkthrough program for the 20GB, 500+ million polygon model of the Boeing 777 aircraft, authoring three invention disclosures relating to that work. He was also one of the principal architects of FlyThru™ Core and CAD Core. He is currently the Principal Investigator for the Access Manager Integration Framework and a new parametric solid modeler. Before coming to Boeing, Eric had worked in computer graphics and CAD for Silicon Graphics Inc., GRAFTEK, the Rensselaer Design Research Center (formally the Center for Interactive Computer Graphics) and the Jet Propulsion Laboratory. He holds a BS and MS in mathematics, and a PhD in applied mathematics from Rensselaer Polytechnic Institute.

James Helman

Member Technical Staff
Silicon Graphics Inc.
2011 N Shoreline Blvd Mountain View CA 94043-1389
jmh@sgi.com

Jim Helman works in Silicon Graphics Advanced Graphics Division as a member of the engineering team for IRIS Performer, SGI's real-time graphics toolkit. Before coming to SGI, he was a student in the Applied Physics Department at Stanford University where he worked on a PhD in data visualization. His interests include virtual environments, game design, and keeping large green cars running.

¹Eric Brechner is now at Microsoft Corporation, One Microsoft Way Redmond, WA 98052-6399
ericbrec@microsoft.com

Ned Greene

Apple Computer and U C Santa Cruz
MS301 3J 1 Infinite Loop Cupertino CA 95014
greene@apple.com

Ned Greene is a member of the Advanced Technology Group at Apple Computer in Cupertino California where he conducts research in computer graphics. He holds a masters degree in computer science from New York University and is currently completing a PhD at the University of California at Santa Cruz. From 1980 to 1989 Ned worked at the Computer Graphics Lab at The New York Institute of Technology where he developed software for computer animation and contributed to pioneering animation projects such as *The Works* and *The Magic Egg*. Over the years he has been a frequent contributor to the SIGGRAPH technical program and Electronic Theater.

Jarek Rossignac

Senior Manager
IBM T.J. Watson Research Center
PO Box 704 Yorktown Heights NY 10598
jarek@watson.ibm.com

Jarek Rossignac is Senior Manager of the 3D Graphics and Interactions Department at IBM Research, covering a wide range of activities in CAD graphics and virtual reality. He joined IBM Research in 85 with a PhD in EE on Solid Modeling from the University of Rochester. In 89 he created the Interactive Geometric Modeling group and managed research projects in topological representations and algorithms for CAD and in the interactive design and inspection of complex animated assemblies. He published over 40 papers winning three Best Paper external awards (Computer and Graphics'90 Eurographics 91 and IEEE Computer Graphics and Applications 93). He authored 12 invention disclosures in modeling and graphics and received numerous IBM internal awards (including the Best Paper award from Computer Science and a Research Division Award for the 3D Interaction Accelerator software system which supports graphics and VR interactions with 3D CAD models of industrial complexity i.e. comprising millions of faces). Rossignac lectures at SIGGRAPH and at Eurographics. He is Associate Editor of four professional journals (ACM Transactions on Graphics, Computer Aided Design, The Visual Computer, and Computer Graphics Forum) and Guest Editor for 7 special issues (three in Computer Aided Design, one in the International Journal of Computational Geometry and Applications, two in the IEEE Computer Graphics and Applications, and one in the ACM transactions on Graphics). He refereed several hundred manuscripts and numerous grant proposals, serves in various international conference committees (including Eurographics 91, 93, 94 and SIGGRAPH 93, 94). Rossignac is also co-chairman of the ACM/SIGGRAPH Symposium on Solid Modeling for 91-95, of the Eurographics 94 Workshop on Graphics Hardware. He is co-chairman of the International Program Committees for Eurographics 96 and for CAD/Graphics 95. Finally, He serves as Program Director for SLAM's Activity Group on Geometric Design.

Thomas A. Funkhouser
MTS
AT&T Bell Laboratories
600 Mountain Avenue, 2A-202 Murray Hill, NJ 07974
funk@research.att.com

Thomas Funkhouser is a member of the technical staff at AT&T Bell Laboratories. His research interests include multi-user systems, global illumination, and algorithms for managing large amounts of three-dimensional data in interactive computer graphics and visualization systems. He is a principal developer of the UC Berkeley Architectural Walkthrough System, which is able to maintain thirty frames per second during interactive visualization of a building model containing 1.5 million polygons. He received a B.S. in biological sciences from Stanford University in 1983, a M.S. in computer science from UCLA in 1989, and a Ph.D. in computer science from UC Berkeley in 1993.