1 Course Title

Out-Of-Core Algorithms for Scientific Visualization and Computer Graphics.

2 Course Organizer

Cláudio T. Silva, AT&T Labs.

3 Course Level

Intermediate — The course is intended for those who have an understanding of the basics of 3D graphics, scientific visualization, analysis of algorithms. The goal of the course is to introduce students and professionals with techniques for handling very large datasets, in particular those datasets which are too large to fit in main memory. The course starts with the basics of external memory algorithms, then covers recently developed out-of-core techniques for a large variety of problems, such as isosurface computation, surface simplification, volume rendering, surface reconstruction, and view-dependent rendering.

4 Proposed Length

Half-day course. (This course can easily be expanded to a full day course if the tutorial reviewers and chairs would find it appropriate).

5 Summary Statement

This course provides an in depth look at external memory algorithms for computer graphics and scientic visualization, and gives insight into developing algorithms for handling very large datasets. The main focus is on core algorithmic ideas, and such knowledge has wide applications in graphics and visualization.

6 Expanded Statement

Input/Output (I/O) communication between fast internal memory and slower external memory is a major bottleneck in many large-scale applications. Algorithms specifically designed to reduce the I/O bottleneck are called *external-memory* algorithms.

This course will focus on describing techniques for handling datasets larger than main memory in scientific visualization and computer graphics. Recently, several external memory techniques have been developed for a wide variety of graphics and visualization problems, including surface simplification, volume rendering, isosurface generation, ray tracing, surface reconstruction, and so on. This work has had significant impact given that in recent years there has been a rapid increase in the raw size of datasets. Several technological trends are contributing to this, such as the development of high-resolution 3D scanners, and the need to visualize ASCI-size (Accelerated Strategic Computing Initiative) datasets. Another important push for this kind of technology is the growing speed gap between main memory and caches, such a gap penalizes algorithms which do not optimize for coherence of access. Because of these reasons, much research in computer graphics focuses on developing out-of-core (and often cache-friendly) techniques.

This course reviews fundamental issues, current problems, and unresolved solutions, and presents an in-depth study of external memory algorithms developed in recent years. Its goal is to provide students and graphics professionals with an effective knowledge of current techniques, as well as the foundation to develop novel techniques on their own.

Here is a short list of representative technical work we will cover in this course:

- "Interactive Out-Of-Core Isosurface Extraction", Y.-J. Chiang, C. Silva, and W. Schroeder, pp. 167–174, IEEE Visualization, 1998.
- "External memory view-dependent simplification", J. El-Sana and Y.-J. Chiang, *Computer Graphics Forum*, 19(3), August 2000.
- "Out-of-core rendering of large unstructured grids", R. Farias and C. Silva, *IEEE Computer Graphics & Applications*, 2001.
- "Out-of-core simplification of large polygonal models", P. Lindstrom, ACM SIGGRAPH 2000, pages 259–262, July 2000.
- "A Memory Insensitive Technique for Large Model Simplification", P. Lindstrom and C. Silva, pp. 121–126, IEEE Visualization 2001.
- "Visualization of large terrains made easy", P. Lindstrom and V. Pascucci, pp. 363–370, IEEE Visualization 2001.

7 Prerequisites

Understanding of the basics of 3D graphics and scientific visualization techniques.

8 Topics List

Basic theory of external memory techniques (e.g., disk modeling, B-trees); isosurface computation; surface simplification; volume rendering; surface reconstruction; coherent ray tracing; and view-dependent rendering.

9 Speakers

Yi-Jen Chiang, Jihad El-Sana, Peter Lindstrom, Cláudio T. Silva.

10 Course Syllabus

- Introduction to the course (Silva)
  5 minutes
- An overview of external memory algorithms (Chiang)
  40 minutes
  Basic theory of external memory techniques (e.g., disk modeling, B-trees); Developing optimal out-of-core algorithms

– Out-of-core scientific visualization (Silva)
  40 minutes
  I/O efficient isosurface computation; volume rendering; streamline computation; multi-resolution methods

– Out-of-core surface simplification (Lindstrom)
  40 minutes

– Out-of-core polygonal rendering (El-Sana)
  40 minutes
  Architectural walkthrough systems; view-dependent rendering; terrain rendering

– External-memory global illumination (Silva)
  30 minutes

– Questions and final remarks (All)
  15 minutes

11 Course History

This is a new course, with mostly all new material. Because large datasets happen quite often in computer graphics and scientific visualization, in the past, some IEEE Visualization (and ACM SIGGRAPH) courses have presented specific approaches for dealing with large datasets in particular domains. In all cases, the main topic of these courses was not external memory computations per se. Instead, external memory issues was often a minor topic, usually not covered in much detail. One of the reasons for such weak coverage of this important subject is that several breakthrough research on the topic is so new. A substantial portion of this important research has been accomplished by the speakers in the course (or their co-workers). Given the importance of the subject, we believe a comprehensive course such as ours would be a significant addition to the course program.

12 Continuity Statement

Chiang starts the course with an introduction to the theory and practice of external memory algorithms. Such a framework is important while designing new algorithms, and assessing the optimality of a particular solution. After that, Silva describes external memory algorithms for several important scientific visualization problems, including isosurface generation, and volume rendering. Lindstrom follows that up with a description of algorithms for surface simplification, and for building level-of-detail hierarchies on disk. This is followed by El-Sana, who describes how to use such data structures for efficient rendering, including rendering data over networks, on remote disks. Finally, Silva covers global illumination algorithms which can be used to generate high-quality images of very large datasets.

13 Course Presenters' Biographies

Cláudio T. Silva
AT&T Labs-Research
180 Park Avenue – Room D265
Florham Park, NJ 07932

E-mail: csilva@research.att.com

Cláudio T. Silva is a Senior Member of Technical Staff in the Information Visualization Research Department at AT&T Labs-Research, and an adjunct assistant professor at the State University of New York at Stony Brook. His current research is on architectures and algorithms for building scalable displays, rendering techniques for large datasets, 3D scanning, and algorithms for graphics hardware. Before joining AT&T, Claudio was a Research Staff Member at the graphics group at IBM T. J. Watson Research Center. Claudio has a Bachelor's degree in mathematics from the Federal University of Ceara (Brazil), and MS and PhD degrees in computer science from the State University of New York at Stony Brook. While a student, and later as an NSF post-doc, he worked at Sandia National Labs, where he developed large-scale scientific visualization algorithms and tools for handling massive datasets. His main research interests are in graphics, visualization, applied computational geometry, and high-performance computing. Claudio has published over 40 papers in international conferences and journals, and presented courses at ACM SIGGRAPH, Eurographics, and IEEE Visualization conferences. He is a member of the ACM, Eurographics, and IEEE. He is on the editorial board of the IEEE Transactions on Visualization and Computer Graphics.

Yi-Jen Chiang
Dept. of Computer and Information Science
Polytechnic University
5 MetroTech Center
Brooklyn, NY 11201-3840
E-mail: yjc@poly.edu

Yi-Jen Chiang is an Assistant Professor of the Computer and Information Science Department at Polytechnic University, Brooklyn New York. His current research focuses on out-of-core algorithms and data structures for large graphics and visualization problems. He developed a series of the first out-of-core techniques for isosurface extraction from large unstructured volume data. He received the Best Paper Award of Eurographics 2000 for his work on external-memory view-dependent simplification and rendering for large polygonal models. Before joining Polytechnic University, he was a post-doc in industrial applied mathematics at SUNY Stony Brook. He has a Bachelor's degree in Computer Science from National Taiwan University (Taiwan), and Sc.M. and Ph.D. degrees in Computer Science from Brown University. His main research interests include graphics, visualization, algorithms and data structures, computational geometry, and high-performance computing. He has published over 20 papers in international journals and conferences in these areas, and received an NSF CAREER award. He is a member of ACM, IEEE, and SIAM.

Jihad El-Sana
Department of Computer Science
Ben Gurion University of The Negeve
Beer-Sheva, 84105
Israel
E-mail: el-sana@cs.bgu.ac.il

Jihad El-Sana is a Lecturer (equivalent to assistant professor in the U.S.) of Computer Science at Ben-Gurion University of the Negev, Israel. El-Sana's research interests include 3D interactive graphics, multi-resolution hierarchies, geometric modeling, computational geometry, virtual environments, and distributed and scientific visualization. His research focus is on polygonal simplification, occlusion culling, accelerating rendering, remote/distributed visualization, and exploring the applications of virtual reality in engineering, science, and medicine. El-Sana received a B.Sc. and M.Sc. in Computer Science from Ben-Gurion University of the Negev, Israel in 1991 and 1993. He received a Ph.D. in Computer Science from the State University of New York at Stony Brook in 1999. Before his Ph.D. studies, El-Sana worked as a scientific programmer at Scitex Corporation to develop commercial image processing and pre-press systems. While a Ph.D. student, he worked at IBM T. J. Watson Research Center on occlusion culling algorithms. El-Sana has published over 20 papers in international conferences and journals. He is a member of the ACM, Eurographics, and IEEE.

Peter Lindstrom
Lawrence Livermore National Laboratory
7000 East Avenue, L-560
Livermore, CA 94551
E-mail: pl@llnl.gov

Peter Lindstrom is a Computer Scientist in the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory (LLNL). His current research is focused on surface simplification and compression, view-dependent rendering and other level-of-detail techniques, and geometric and multiresolution modeling. As a member of the ASCI VIEWS visualization team at LLNL, he works on out-of-core methods for managing and visualizing terascale geometric data sets generated in large scientific simulations. Peter joined LLNL in 2000 after receiving his PhD in computer science from the Georgia Institute of Technology, and graduated with BS degrees in computer science, mathematics, and physics from Elon University, North Carolina, in 1994. Peter has published and presented papers at ACM SIGGRAPH, ACM TOG, IEEE Visualization, and IEEE TVCG, among other computer graphics conferences and journals. He is a member of ACM and IEEE.

14 Special Notes

The course notes are going to include: printed slides of the presentations; additional text notes (to accompany slides); and selected re-prints of relevant papers on external memory algorithms.

For our presentation, we require the basic course configuration (PC with fast 3D acceleration).