

---

ACM SIGGRAPH Asia 2013 Course Notes

# GPU-Based Large-Scale Visualization

---

**Markus Hadwiger**

KAUST, Saudi Arabia

**Jens Krüger**

University of Duisburg-Essen, Germany

**Johanna Beyer**

Harvard University, USA

**Stefan Bruckner**

University of Bergen, Norway



**SIGGRAPH**  
**ASIA 2013**  
**HONG KONG**

---

# Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Abstract . . . . .	1
1.2	Level of Difficulty . . . . .	2
1.3	Intended Audience . . . . .	2
1.4	Course Prerequisites . . . . .	2
1.5	Online Course Material . . . . .	2
<b>2</b>	<b>Syllabus</b>	<b>3</b>
<b>3</b>	<b>Speaker Biographies</b>	<b>6</b>

# Introduction

---

An overview of modern techniques for interactive and scalable visualization of large-scale image and volume data on GPUs via display-aware multi-resolution algorithms and on-demand streaming of data. We describe recent GPU data structures and advanced multi-resolution representations, virtual texturing approaches, and out-of-core techniques.

## 1.1 Abstract

Recent advances in image and volume acquisition as well as computational advances in simulation have led to an explosion of the amount of data that must be visualized and analyzed. Modern techniques combine the parallel processing power of GPUs with out-of-core methods and data streaming to enable the interactive visualization of giga- and terabytes of image and volume data. A major enabler for interactivity is making both the computational and the visualization effort proportional to the amount of data that is actually visible on screen, decoupling it from the full data size. This leads to powerful display-aware multi-resolution techniques that enable the visualization of data of almost arbitrary size.

The course consists of two major parts: An introductory part that progresses from fundamentals to modern techniques, and a more advanced part that discusses details of ray-guided volume rendering, novel data structures for display-aware visualization and processing, and the remote visualization of large online data collections.

You will learn how to develop efficient GPU data structures and large-scale visualizations, implement out-of-core strategies and concepts such as virtual texturing that have only been employed recently, as well as how to use modern multi-resolution representations. These approaches reduce the GPU memory requirements of extremely large data to a working set size that fits into current GPUs. You will learn how to perform ray-casting of volume data of almost arbitrary size and how to render and process gigapixel images using scalable, display-aware techniques. We will describe custom virtual texturing architectures as well as recent

hardware developments in this area. We will also describe client/server systems for distributed visualization, on-demand data processing and streaming, and remote visualization.

We will describe implementations using OpenGL as well as CUDA, exploiting parallelism on GPUs combined with additional asynchronous processing and data streaming on CPUs.

## 1.2 Level of Difficulty

Intermediate

## 1.3 Intended Audience

We target researchers and practitioners in visualization and computer graphics that want to learn recent GPU techniques and hardware/API capabilities for implementing visualization systems that scale to very large data. The course will also be interesting for people interested in processing and rendering gigapixel images in a scalable manner.

## 1.4 Course Prerequisites

Course participants should have a basic understanding of how GPUs work and how they are programmed using standard graphics APIs such as OpenGL or Direct3D, and a basic understanding of general-purpose GPU programming using CUDA or OpenCL.

## 1.5 Online Course Material

Further material for this course can be found online at:

<http://people.seas.harvard.edu/~jbeyer/LargeScaleVisCourse/>

# Syllabus

---

## Introduction, Basics of Scalable Visualization and GPU Virtual Memory Architectures

**Markus Hadwiger** [60 min]

After introducing the course, I will describe GPU-based visualization techniques, focusing on GPU ray-casting for large volume data. We will look at fundamental data structures such as grids and trees, and discuss how to perform traversal on GPUs. We will discuss the basics of rendering large images and volumes and out-of-core approaches, and I will describe the basic ray-casting optimization techniques of empty space skipping and early ray termination with a focus on large data. I will give an overview of virtual texturing on GPUs and how virtual memory hierarchies can be employed for the interactive visualization of data of almost arbitrary size.

## Scalable Volume Visualization Architectures and Applications

**Johanna Beyer** [45 min]

I will start by explaining some advanced techniques for large-scale ray-casting, including multi-level page tables for accessing terabyte-sized volumes on the GPU, and demonstrate modern empty space skipping techniques for increasing rendering performance. In the second part of my talk, I will introduce and compare different volume processing architectures for handling large data. Data streaming and on-demand construction of 3D data blocks allows to circumvent expensive pre-processing and allows to process and stream data entirely on-demand, thereby limiting the amount of data that has to be handled to the visible subset on the screen. In the last part of my talk I will present higher-level applications

that are built on top of current GPU-based visualization systems for large-scale data. I will present a system for query-guided visual analysis of large volumetric electron microscopy data, built on top of a powerful petascale visualization system.

**Break [15 min]**

## GPU-Based Ray-Guided Volume Rendering Algorithms

**Jens Krüger [40 min]**

GPU-Based Ray-Guided Volume Rendering algorithms have been demonstrated to be the most efficient approaches for interactive volume rendering of large data. For such data sets, the data movement strategies from background memory such as disk or network to main memory and ultimately to GPU memory play an important role. In this talk I will take a closer look at data storage and layout strategies. I will detail a number of design aspects in ray-guided volume rendering, demonstrating that a whole-system design viewpoint achieves superior results as compared to individual optimizations.

## Display-Aware Visualization and Processing, and PDF-Based Multi-Resolution Representations

**Markus Hadwiger [40 min]**

I will give an overview of display-aware techniques for the visualization and processing of gigapixel images and large-scale volume data, where the computational effort is proportional to the actually visible data instead of to the original data size. I will describe a novel probabilistic data structure called sparse pdf maps/volumes that encodes the information of high-resolution data in a multi-resolution pyramid using probability density functions (pdfs) of local pixel or voxel neighborhoods. This enables the accurate visualization of high-resolution data with arbitrary transfer functions at any output resolution. Sparse pdf maps/volumes also enable image processing operations such as local Laplacian filters, smoothed local histogram filters, and bilateral filters to be computed in a display-aware, and thus scalable, manner.

## Remote and Web-based Visualization of Data Collections

**Stefan Bruckner** [40 min]

I will present approaches for the visualization of large data collections using GPU-based rendering. Data repositories in disciplines such as medicine or biology grow in number and size which means that the cost of data transfer frequently becomes prohibitive and image-based solutions are the only viable alternative. The talk will provide an overview of software architectures and solutions for such scenarios. In particular, I will focus on approaches which employ web technology to enable interactive remote rendering and show how they can be combined with client-side interfaces to provide high-fidelity visualization applications.

# Speaker Biographies

---

## **Markus Hadwiger**

**Assistant Professor**

**King Abdullah University of Science and Technology  
Saudi Arabia**

`markus.hadwiger@kaust.edu.sa`

`http://www.kaust.edu.sa/markushadwiger`

Markus Hadwiger is an Assistant Professor in computer science at King Abdullah University of Science and Technology (KAUST) and head of the High-Performance Visualization group at the Geometric Modeling and Scientific Visualization Center at KAUST. Before joining KAUST, he was a Senior Researcher at the VRVis Research Center in Vienna. He received a PhD in computer science from the Vienna University of Technology in 2004. He is a co-author of the book Real-Time Volume Graphics. His research interests include petascale visual computing and scientific visualization, volume rendering, large-scale image processing, and GPU algorithms and architecture.

## **Jens Krüger**

**Professor**

**University of Duisburg-Essen  
Germany**

`jens.krueger@uni-due.de`

`http://hpc.uni-due.de/krueger`

Jens Krüger finished his PhD in 2006 at the Technische Universität München. After Post Doc positions in Munich and at the Scientific Computing and Imaging (SCI) Institute he became Research Assistant Professor at the University of Utah. In 2009 he joined the Cluster of Excellence to head the Interactive Visualization and Data Analysis group. Since 2013 he is Chair of the High Performance Computing group at the

University of Duisburg-Essen. He also holds an adjunct professor title of the University of Utah and is a principal investigator in the Intel Visual Computing Institute.

## **Johanna Beyer**

**Postdoctoral Fellow**  
**Harvard University**  
**USA**

[jbeyer@seas.harvard.edu](mailto:jbeyer@seas.harvard.edu)  
<http://sites.google.com/site/johannabeyercv>

Johanna Beyer is a postdoctoral fellow in the School of Engineering and Applied Sciences at Harvard University. Before joining Harvard, she was a postdoctoral fellow at the Geometric Modeling and Scientific Visualization Center at King Abdullah University of Science and Technology. Her research interests include large-data visualization, parallel visualization, and GPU-based volume rendering for neuroscience and neurobiology. She received a PhD in computer science from the Vienna University of Technology in 2010.

## **Stefan Bruckner**

**Professor**  
**University of Bergen**  
**Norway**

[stefan.bruckner@uib.no](mailto:stefan.bruckner@uib.no)  
<http://www.ii.uib.no/vis/team/bruckner>

Stefan Bruckner received his Ph.D. in 2008 from Vienna University of Technology. He was awarded the habilitation (*venia docendi*) in Practical Computer Science in 2012. From 2008 to 2013, he was an assistant professor at the Institute of Computer Graphics and Algorithms at VUT. In 2009/2010, he spent one year as a visiting Postdoctoral Research Fellow at Simon Fraser University, Canada. Since March 2013, he is a full professor in visualization at the University of Bergen. His research interests include biomedical and illustrative visualization, volume rendering, and visual data exploration. In 2011, he was recipient of the Eurographics Young Researcher Award.

