Welcome to a day with talks on

Accelerating Computations

Research conference on graphics processing units, visual computing and beyond...

On 15 December 2011 the Computer Graphics Lab at the Alexandra Institute in Aarhus will host a conference on the core aspects of accelerating computations. This research conference will be divided into a technical part called Technical Talks in the morning and a business-oriented part called Applications Talks in the afternoon.

The conference is free and you can participate either in the Technical Talks, the Applications Talks or both. Keynote speaker will be David McAllister from NVIDIA.

The Computer Graphics Lab at the Alexandra Institute is working on realising efficient computations on massive data sets by using modern many-core processors and massive data algorithms. Subsequently interactive information visualisation is crucial to analyse and understand massive data sets.

PROGRAMME

10:00 - 10:15 **Introduction**

10:15 – 14:45 Technical Talks

10:15 – 11:00 **How OptiX Makes the GPU Shine – a look inside NVIDIA's Ray Tracing Engine**
David McAllister, NVIDIA

11:00 – 11:45 **MR reconstruction on GPU**
Thomas Sangild, Computer Science, Aarhus University

11:45 – 12:30 **Lunch**

12:30 - 13:15 **Parthenon Renderer Revealed**
Toshiya Hachisuka, Computer Science, Aarhus University

13:15 - 13:45 **Subsurface Light Propagation Volumes**
Thomas Kim Kjeldsen, Computer Graphics Lab, Alexandra Institute

13:45 – 14:00 **Coffee break**

14:00 – 14:45 **Accelerating Dense Linear Algebra on the GPU**
Hans Henrik Brandenborg Sørensen, DTU Informatics

14:45 – 16:45 Applications Talks

14:45 – 15:15 **Massive acceleration**
Jesper Mosegaard, Computer Graphics Lab, Alexandra Institute

15:15 - 15:30 **Coffee break**

15:30 – 16:15 **Dozens of Uses for Billions of Rays – a survey of ray tracing applications**
David McAllister, NVIDIA

16:15 – 16:45 **Scalable GPU Computing Service Architecture: LEGO 3DServices**
Henrik Høj Madsen, LEGO, and Michael Schøler, Hinnerup.net

16:45 – **Networking and sandwiches**

Venue: Alexandra Institute, Peter Bøgh Andersen Auditorium, Nygaard building 5335, Finlandsgade 21-23, 8200 Aarhus N

Registration: Online at http://cg.alexandra.dk/signup/ – no later than 12 December

NVIDIA LEGO Hinnerup Net DTU ALEXANDRA INSTITUTTET

How OptiX Makes the GPU Shine – a look inside NVIDIA's Ray Tracing Engine

David McAllister, Optix Manager, NVIDIA
I will briefly describe the OptiX programming model, then dive into the internals of how we exposed the GPU's computational power for ray tracing in an application programmable way.

MR reconstruction on GPU

Thomas Sangild Sørensen, Associate Professor, Computer Science, Aarhus University
A barrier to the adoption of non-Cartesian parallel magnetic resonance imaging for real-time applications has been the times required for the image reconstructions. These times have exceeded the underlying acquisition time thus preventing real-time display of the acquired images. We present a reconstruction algorithm for commodity graphics hardware (GPUs) to enable real time reconstruction of sensitivity encoded radial imaging (radial SENSE).

Parthenon Renderer Revealed

Toshiya Hachisuka, Assistant Professor, Computer Science, Aarhus University
Parthenon Renderer, initially released back in 2002, is one of the earliest publicly available rendering software that utilise graphics hardware for accelerating computation of high-quality offline rendering. I will talk about the inside of Parthenon Renderer in order to give you some examples of engineering choices and algorithm design that make (and made) sense for an offline rendering system using graphics hardware.

Subsurface Light Propagation Volumes

Thomas Kim Kjeldsen, Research and Innovation Scientist, Computer Graphics Lab, Alexandra Institute
We present the Subsurface Light Propagation Volume (SSLPV) method for real-time approximation of subsurface scattering effects in dynamic scenes with changing mesh topology and lighting. SSLPV extends the Light Propagation Volume (LPV) technique for indirect illumination in video games. We introduce a new consistent method for injecting flux from point light sources into an LPV grid, a new rendering method which consistently converts light intensity stored in an LPV grid into incident radiance, as well as a model for light scattering and absorption inside heterogeneous materials. Our scheme does not require any precomputation and handles arbitrarily deforming meshes. We show that SSLPV provides visually pleasing results in real-time at the expense of a few milliseconds of added rendering time.

Accelerating Dense Linear Algebra on the GPU

Hans Henrik Brandenborg Sørensen, Post. Doc., GPU Lab, DTU Informatics
GPUs have already become an integral part of high performance scientific computing, since they offer dedicated parallel hardware that can potentially accelerate the execution of many scientific applications. In this talk, I will consider the automatic performance acceleration of dense vector and matrix-vector operations on GPUs. Such operations form the backbone of level 1 and level 2 routines in the Basic Linear Algebra Subroutines (BLAS) library and are therefore of great importance in many scientific applications.

The target hardware is the most recent NVIDIA Tesla 20-series (Fermi architecture). Most of the techniques I discuss for accelerating dense linear algebra are applicable to memory-bound GPU algorithms in general.

Massive Acceleration at the Alexandra CG Lab

Jesper Bjerg Mosegaard, Head of Research and Innovation, Computer Graphics Lab, Alexandra Institute
The Computer Graphics Lab at the Alexandra Institute does research and development within the topic of fast and accurate simulation and visualisation in high quality. This talk describes the role of the Alexandra Institute in transferring research to application in Danish Industry as well as specific opportunities for companies to benefit from the latest knowledge and technology.

Dozens of Uses for Billions of Rays – a survey of ray tracing applications

David McAllister, Optix Manager, NVIDIA
Since introducing OptiX in 2009, NVIDIA has been approached by engineers from industries as diverse as geothermal exploration, cell phone antenna design, and automotive headlamp design that have one thing in common – the need to intersect rays, usually billions of them, against a database of geometry. I will survey many applications of ray tracing, within and beyond computer graphics, and show how accelerating ray tracing using GPUs addresses many challenges in industry.

Scalable GPU Computing Service Architecture: LEGO 3DServices

Michael Schøler, Hinnerup.net and
Henrik Høj Madsen, Solution Architect, LEGO
As LEGO is moving into the virtual playspace, a platform technology has been developed in-house primarily based on NVIDIA technologies, featuring:

- CUDA, OptiX, OpenGL and general shaders
- 17 Quadro Plex in multiple environments, multiple datacentres
- Advanced shading techniques for approaching high-quality results in real-time
- On-demand asset generation
- CDN assets distribution
- A generic service-oriented interface
- A distributed rendering architecture
- Architectural patterns for distributed computing
- A plugin architecture supporting existing and future LEGO experiences
- A mentality shift from traditional ways of doing things on CPU vs GPU
- General experiences from developing on NVIDIA tech in a large-scale Enterprise setup

The LEGO 3DServices system is designed to support diverse computational needs such as on-demand rendering, mesh optimisation, a Massive Multiplayer Online Game (MMO), product visualisations, 3D modeling and other current and future demanding computational tasks.

Our aim with this session is to share our learnings and present LEGO's vision of the future of distributed GPU accelerated computation as a business-driven platform technology.