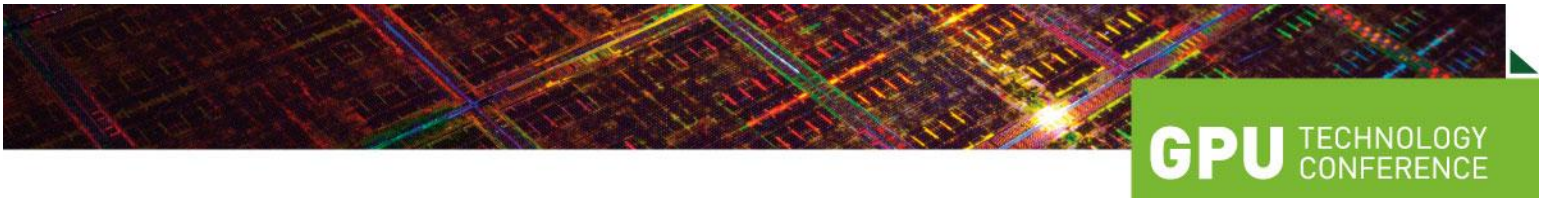

GPU Technology Conference, May 14-17, 2012
McEnery Convention Center, San Jose, California
www.gputechconf.com

Sessions on Cloud Computing (subject to change)

IMPORTANT: Visit http://www.gputechconf.com/page/sessions.html for the most up-to-date schedule.

S0337 - High-Throughput Epistasis Screening Using GPUs
Mark Seligman (Insilicos LLC)
Day: Tuesday, 05/15 | **Time:** 9:00 am - 9:25 am
Topic Areas: Bioinformatics; Life Sciences; Supercomputing; Cloud Computing
Session Level: Intermediate

Epistasis is the interaction of two or more genes in coding for a biological property. Epistasis is believed to be an important factor in an individual's susceptibility to disease, and the search for epistasis is a major component in the development of personalized approaches to genomic medicine. Statistical tests for epistasis are typically confounded by the multiple-testing problem, that is, the aggregated loss of precision incurred through repeated hypothesis testing. One way to circumvent this problem is to simulate a false-discovery rate via resampling. We report success in using GPUs to accelerate these highly compute-intensive resampling techniques.

S0254 - Graphics in the Cloud - How NVIDIA is Enabling Cloud Visualization
Will Wade (NVIDIA)
Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm
Topic Areas: Cloud Computing; Visualization; Computer Graphics
Session Level: Intermediate

Engineers, artists, scientists, and gamers are the most demanding visual thinkers on the planet, and as such have not been willing to move their computing environments to the infamous "cloud". These remotely accessed systems are seen as slow and not up to the visual experience that users expect when dealing with these types of applications. NVIDIA aims to change that perception with the NVIDIA Virtual Graphics Platform. In this session you will hear about the technologies behind accelerating graphics in the cloud, and some of the industry partnerships that are enabling it.

S0413 - Delivering 3D Professional Graphics from the Cloud with Citrix XenDesktop
Derek Thorslund (Citrix Systems, Inc.)
Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Cloud Computing; Computer Graphics; Visualization
Session Level: Beginner

Recent technological advances have made it practical to deliver 3D professional graphics applications from the Cloud (private or public) with a high quality user experience and at an attractive cost. Organizations can keep their intellectual property safe in the data center since only fully-rendered screen images are sent over the network. Users in remote locations no longer have to wait for large file transfers. And they can access 3D models from a wide variety of devices, including iPads and Android tablets. Learn how Citrix XenDesktop, XenServer and Receiver technologies have made all of this a reality for many organizations today.

S0261 - Scalable GPU Computing Service Architecture
Henrik Høj Madsen (LEGO), Michael Schøler (LEGO)
Day: Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm
Topic Areas: Cloud Computing; Computer Graphics; Ray Tracing
Session Level: Intermediate

In this session we describe our GPU accelerated computing service which supports several internal business processes in a large scale company setup. The service supports diverse computational needs such as on-demand rendering, mesh optimization, a Massive Multiplayer Online Game (MMO), product visualizations and other demanding computational tasks. We present the architectural considerations for a service-oriented computational framework and the practical learning's and opportunities encountered during development a enterprise system using NVIDIA technologies such as CUDA, OptiX, OpenGL and OpenCL. Our aim is to share knowledge and present LEGO's vision for a GPU accelerated computational platform as a business-driven technology.

S0359 - VMware and NVIDIA: Delivering 3D Workstations from the Cloud
Aaron Blasius (VMware), Warren Ponder (VMware)
Day: Tuesday, 05/15 | **Time:** 5:00 pm - 5:50 pm
Topic Areas: Visualization; Cloud Computing
Session Level: Advanced

This session will detail the delivery of the most demanding Workstation class workloads from the private cloud using technologies from NVIDIA and VMware. We will cover the configuration and performance metrics of the combined VMware, NVIDIA direct pass through hardware accelerated graphics solution. Using sample workloads, we will demonstrate how customers can realize the operational and security benefits of cloud based personal computing without sacrificing performance.

S0405 - New Generation GPU Accelerated Financial Quant Libraries
Daniel Egloff (QuantAlea GmbH)
Day: Wednesday, 05/16 | **Time:** 3:00 pm - 3:50 pm
Topic Areas: Finance; Application Design & Porting Techniques; Algorithms & Numerical Techniques; Cloud Computing
Session Level: Advanced

Learn from industry experts how new generation GPU accelerated solutions for derivative pricing, hedging, and risk management can be build more efficiently with modern technology and functional programming languages like F# on .NET or Scala on the Java VM. As a concrete example we report from a large derivative pricing project developed in F# on .NET. We will introduce the key design concepts and parallelization strategies, which lead to an efficient and transparent GPU acceleration. Several examples will illustrate the benefit of the functional as compared to the classical object oriented approach.

S0507 - Interactive and Scalable Subsurface Data Visualization Framework
Tom-Michael Thamm (NVIDIA ARC), Marc Nienhaus (NVIDIA ARC)
Day: Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm
Topic Areas: Visualization; Cloud Computing
Session Level: Intermediate

The goal is to present an interactive visualization framework for large geo-spatial data. This framework has been developed by NVIDIA Advanced Rendering Center for the oil and gas (Hydrocarbone) industry. The CUDA based application is running on the cloud at interactive frame-rates. The visualization is remote on clients in a browser,

including tablets. The scalable visualization framework can handle terra bytes of.

S0411 - Artifact-Free Cloud-Based CAD Rendering
Sara McMains (UC Berkeley), Sushrut Pavanaskar (UC Berkeley)
Day: Thursday, 05/17 | **Time:** 4:30 pm - 4:55 pm
Topic Areas: Algorithms & Numerical Techniques; Computer Graphics; Cloud Computing; Visualization
Session Level: Beginner

Cloud computing for mechanical CAD provides centrally stored and synchronized models for concurrent engineering. For compactness, trimmed parametric NURBS surface representations are optimal for data transfer to client devices, which must evaluate and render models locally. Direct GPU rendering without pre-tessellation is an attractive solution in this context, both for speed and to preserve fidelity to the original geometry. However, existing data-parallel direct rendering approaches for NURBS suffer from rendering artifacts at trim boundaries. This talk proposes a solution to address these rendering artifacts that are still preventing wide-scale adoption of all such direct rendering algorithms for trimmed parametric models.