

GPU Technology Conference 2010 Sessions on Databases & Data Mining (subject to change)

IMPORTANT: Visit www.nvidia.com/qtc for the most up-to-date schedule and to enroll into sessions to ensure your spot in the most popular courses.

2237 - Accelerating Business Intelligence Applications with Fast Multidimensional Aggregation

In this research session, we present an approach using NVIDIA GPUs as massively parallel coprocessors for in-memory OLAP computations. Early tests have shown speedup factors of more than 40x compared to optimized sequential algorithms on a CPU. In addition to the data structures and algorithms involved, we describe a method to extend the approach to systems with more than one GPU in order to scale it to larger data sets.

Speakers: Tobias Lauer, University of Freiburg, Christoffer Anselm, Jedox AG

Topic: Databases & Data Mining

Time: Wednesday, September, 22nd, 15:00 - 15:50

4004 - Emerging Companies: CEO on Stage featuring Cooliris, Jedox Business Intelligence, and Playcast

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature Cooliris, Empulse, and Playcast - covering the fields of computer graphics, video processing and data mining.

Panelists will include Drew Lanza (Partner, Morgenthaler), Charles Carmel (Cisco), and Nathan Brookwood (Principal Analyst, Insight64).

Speakers: Drew Lanza, Morgenthaler, Natan Peterfreund, Playcast Media Systems, Nathan Brookwood, Insight 64, Charles Carmel, Cisco, Michael Hummel, empulse GmbH

Topics: General Interest, Databases & Data Mining, Video Processing, Computer Graphics

Time: Wednesday, September, 22nd, 16:00 - 16:50

4005 - Emerging Companies: CEO on Stage featuring Empulse, Rocketick, and Softkinetic

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature Jedox Business Intelligence, Rocketick, and Softkinetic - covering the fields of computer vision, data mining and high performance computing.

Panelists will include Drew Lanza (Partner, Morgenthaler), Charles Carmel (Cisco), and Nathan Brookwood (Principal Analyst, Insight64).

Speakers: Drew Lanza, Morgenthaler, Uri Tal, Rocketick, Charles Carmel, Cisco, Nathan Brookwood, Insight 64, Kristian Raue, Jedox AG, Michel Tombroff, Softkinetic

Topics: General Interest, Computer Vision, Databases & Data Mining, High Performance Computing

Time: Wednesday, September, 22nd, 17:00 - 17:50

2069 - GPU-Accelerated Business Intelligence Analytics

Join us and learn why GPU computing is a game changer for business intelligence (BI). We will discuss how GPUs can be used to accelerate BI analytics at much lower cost, higher performance, and better power efficiency than other alternatives.

Speaker: Ren Wu, HP Labs

Topics: Databases & Data Mining, Finance, High Performance Computing

Time: Wednesday, September, 22nd, 16:00 - 16:50

2120 - High Performance Complex Event Processing on GPGPU

Complex Event processing (CEP), a crucial component in enterprise-scale applications, is the key element in that it allows applications to process the incoming event streams and apply relevant techniques in real-time for quicker decisions, making it easy to identify complex patterns in the

events. Much of the time, this system is consumed by the event matching algorithms. Our work utilizes the highly parallel GPU for event matching algorithm wherein every incoming event is worked upon by this algorithm and results in high throughput.

Speakers: Murali Krishna, Infosys Technologies Limited, Dr. Sudeep, Infosys

Topics: Databases & Data Mining, Finance

Time: Wednesday, September, 22nd, 14:00 - 14:50

2092 - Integrating CUDA into a Large-Scale Commercial Database Management System

In a large-scale database installation where data tables are distributed across multiple servers, computational throughput can be optimized by using GPUs on each server and integrating database management with GPU resources. In the Department of Physics and Astronomy at The Johns Hopkins University, we are experimenting with a set of software tools that closely couple SQL statements with GPU functionality. While still under development, the new framework is now routinely used in our research projects, e.g., to study the spatial clustering of galaxies as well as genomics.

Speakers: Richard Wilton, The Johns Hopkins University, Tamas Budavari, Johns Hopkins University, Alex Szalay, The Johns Hopkins University

Topics: Databases & Data Mining, Astronomy & Astrophysics, High Performance Computing, Tools & Libraries

Time: Wednesday, September, 22nd, 11:00 - 11:50

2014 - Scalable Subsurface Data Visualization Framework

Mental Images' DiCE-based geospatial library is a CUDA and cluster-based visualization framework that enables scalable processing and rendering of huge amounts of subsurface data for interactive seismic interpretation.

Geospatial exploration in the oil and gas industries is concerned with scanning the earth's subsurface structure for detecting oil and for cost-effective drilling of detected oil reservoirs.

Efficient seismic interpretation requires the interpreters to be able to interactively explore huge amounts of volumetric seismic information with embedded stacked horizons to gain visual insight into the subsurface structure and to determine where oil recovery facilities and drilling infrastructure shall be built.

Speakers: Tom-Michael Thamm, mental images GmbH, Marc Nienhaus, mental images GmbH

Topics: Energy Exploration, Databases & Data Mining, Imaging, Tools & Libraries

Time: Wednesday, September, 22nd, 17:00 - 17:50

2140 - Superfast Nearest Neighbor Searches Using a Minimal kd-tree

Learn how to adapt a kd-tree spatial data structure for efficient nearest neighbor (NN) searches on a GPU. Although the kd-tree is not a natural fit for GPU implementation, it can still be effective with the right engineering decisions. By bounding the maximum height of the kd-tree, minimizing the memory footprint of data structures, and optimizing the GPU kernel code, multi-core GPU NN searches with tens of thousands to tens of millions of points run 10-40 times faster than the equivalent single-core CPU NN searches.

Speaker: Shawn Brown, UNC, Chapel Hill

Topics: Algorithms & Numerical Techniques, Databases & Data Mining, Machine Learning & Artificial Intelligence

Time: Wednesday, September, 22nd, 14:00 - 14:50

2111 - Using R for High-Performance Data Analysis

Data analysis is the art and the science of getting the correct quantitative models and their numerical parameters from the observed data. In this talk, we report on a project to integrate CUDA into the open source data analysis environment R. The combined use of the CPU and GPU resources can efficiently exploit the significant amount of data parallelism inherent in most data analysis problems and methods. This makes interactive analysis possible even for large, compute-intensive problems. The implementation and the achievable performance gains will be demonstrated on a concrete example from quantitative finance.

Speaker: Domokos Vermes, Worcester Polytechnic Institute

Topics: Tools & Libraries, Databases & Data Mining, Finance, Life Sciences

Time: Tuesday, September, 21st, 16:30 - 16:50