GPU Technology Conference 2010 Sessions on Life Sciences

(subject to change)

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

2055 - Application of Fermi GPU to Flow Cytometry and Cancer Detection

Learn how a Tesla C2050 enabled scientists to explore cancer data sets 400 times faster than a PC-only implementation. Discusses how the results of this work may lead to better diagnostics for detecting leukemia in blood cells.

Speaker: Robert Zigon, Beckman Coulter

Topic: Life Sciences

Time: Thursday, September, 23rd, 10:00 - 10:50

2105 - CUDA-FRESCO: An Efficient Algorithm for Mapping Short Reads

Learn about CUDA-FRESCO and how it addresses issues with MUMmerGPU. We will detail how CUDA-FRESCO overcomes MUMmerGPU's problems processing reads with errors or mismatches and delivers additional performance beyond MUMmerGPU's 5-12x speedup with less than 100bp query length.

Speaker: Chun-Yuan Lin, Department of CSIE, Chang Gung University

Topics: Life Sciences, Algorithms & Numerical Techniques, Tools & Libraries

Time: Thursday, September, 23rd, 15:00 - 15:50

1002 - Day 2 Keynote with Dr. Klaus Schulten, University of Illinois at Urbana-Champaign

How does the H1N1 “Swine Flu” virus avoid drugs while attacking our cells? What can we learn about solar energy by studying biological photosynthesis? How do our cells read the genetic code? What comes next in computational biology?

Computational biology is approaching a new and exciting frontier: the ability to simulate structures and processes in living cells. Come learn about the "computational microscope," a new research instrument that scientists can use to simulate biomolecules at nearly infinite resolution. The computational microscope complements the most advanced physical microscopes to guide today's biomedical research. In this keynote address, computational biology pioneer Dr. Klaus Schulten of the University of Illinois, Urbana-Champaign, will introduce the computational microscope, showcase the widely used software underlying it, and highlight major discoveries made with the aid of the computational microscope ranging from viewing protein folding, translating the genetic code in cells, and harvesting solar energy in photosynthesis. He will also look towards a future when cell tomography and computing will establish atom-by-atom views of entire life forms.

Klaus Schulten received his Ph.D. from Harvard University in 1974. He is Swanlund Professor of Physics and is also affiliated with the Department of Chemistry as well as with the Center for Biophysics and Computational Biology. Professor Schulten is a full-time faculty member in the Beckman Institute and directs the Theoretical and Computational Biophysics Group at the University of Illinois Urbana-Champaign, IL.

Honors and awards: Award in Computational Biology 2008; Humboldt Award of the German Humboldt Foundation (2004); University of Illinois Scholar (1996); Fellow of the American Physical Society (1993); Nernst Prize of the Physical Chemistry Society of Germany (1981).

Topics: General Interest, Life Sciences

Time: Wednesday, September, 22nd, 09:00 - 9:50

2007 - Folding@home: Petaflops on the Cheap Today; Exaflops Soon?

Learn how Folding@home has used petascale computing with GPUs to make fundamental breakthroughs in computational biology and how this technology can make an impact in your work.

Speaker: Vijay Pande, Stanford University

Topics: Life Sciences, Cloud Computing, High Performance Computing, Molecular Dynamics

Time: Thursday, September, 23rd, 11:00 - 11:50

2282 - GPU-Enabled Biomedical Imaging

The purpose of this presentation is to describe several novel biomedical imaging applications which make extensive use of GPUs. In CT iterative reconstructions, for example, high

performance computing is allowing us to see details and structures we previously were not able to discern.

Speaker: Homer Pien, MGH / HMS

Topics: Medical Imaging & Visualization, High Performance Computing, Imaging, Life Sciences

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

2062 - HOOMD-blue: Fast and Flexible Many-Particle Dynamics

See the newest capabilities and performance enhancements in HOOMD-blue, a general-purpose many-particle dynamics application written for GPUs. Speedups of 80-100x are attained for a wide range of simulation types. Topics for this presentation include an overview of HOOMD-blue, design and implementation details of the underlying algorithms, and a discussion on how generality is maintained without sacrificing performance.

Speaker: Joshua Anderson, University of Michigan

Topics: Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

Time: Thursday, September, 23rd, 15:00 - 15:50

2073 - High Performance Molecular Simulation, Visualization, and Analysis on GPUs

This talk will present recent successes in the use of GPUs to accelerate interactive visualization and analysis tasks on desktop computers, and batch-mode simulation and analysis jobs on GPU-accelerated HPC clusters. We'll present Fermi-specific algorithms and optimizations and compare with those for other devices. We'll also present performance and performance/watt results for NAMD molecular dynamics simulations and VMD analysis calculations on GPU clusters, and conclude with a discussion of ongoing work and future opportunities for GPU acceleration, particularly as applied to the analysis of petascale simulations of large biomolecular complexes and long simulation timescales.

Speaker: John Stone, University of Illinois at Urbana-Champaign

Topics: Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

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

2030 - High-Throughput Cell Signaling Network Learning with GPUs

Explore how GPUs are being used to enable high-throughput cell signaling network discovery and data-intensive computational systems biology more generally. Systems biology is transitioning from a largely reductive discipline to one focused on building predictive models of large-scale biological systems. New instrumentation will provide the necessary raw data for such an approach, the key challenge now is building the hardware and software tools to efficiently and interactively build these models. This session will describe how GPUs can and will play a key role in these efforts.

Speaker: Michael Linderman, Stanford University

Topics: Life Sciences, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence

Time: Thursday, September, 23rd, 09:00 - 9:50

2139 - Interactive Histology of Large-Scale Biomedical Image Stacks

Get the latest information on leveraging GPU computing to process and visualize large-scale biomedical image stacks. We will discuss both display-aware processing and GPU-accelerated texture compression for histology applications on the GPU.

Speakers: Won-Ki Jeong, Harvard University, Jens Schneider, King Abdullah University of Science and Technology

Topics: Medical Imaging & Visualization, Imaging, Life Sciences

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

2163 - Leveraging GPUs for Evolutionary Game Theory

Learn how GPUs are being used to accelerate the study of the emergence of cooperative behavior in biology, from the interactions of humans to viruses to bacteria. The work presented here achieves a speedup of 209x on a cluster of 4 Tesla GPUs.

Speaker: Amanda Peters, Harvard University

Topics: Algorithms & Numerical Techniques, Life Sciences

Time: Wednesday, September, 22nd, 10:00 - 10:20

2203 - Modeling Evolution Computing the Tree of Life

Learn how GPUs are being used to accelerate our understanding of the tree of life. This session will cover BEAGLE, which is an open API and library for evaluating phylogenetic likelihoods of biomolecular sequence evolution. BEAGLE uses novel algorithms and methods for evaluating phylogenies under arbitrary molecular evolutionary models on GPUs, making use of the large number of processing cores to efficiently parallelize calculations.

Speaker: Daniel Ayres, University of Maryland

Topic: Life Sciences

Time: Thursday, September, 23rd, 11:00 - 11:50

2115 - Modified Smith-Waterman-Gotoh Algorithm for CUDA Implementation

It is axiomatic that computational throughput can be increased by exploiting the parallelism of GPU hardware –- but what if the computational algorithm is not easy to implement in parallel? We have modified one such algorithm -– the Smith-Waterman-Gotoh dynamic programming algorithm for local sequence alignment -– so as to make it more amenable to data-parallel computation. The result is a successful CUDA implementation that fully exploits GPU parallelism.

Speaker: Richard Wilton, The Johns Hopkins University

Topics: Life Sciences, Algorithms & Numerical Techniques

Time: Thursday, September, 23rd, 14:00 - 14:50

2054 - NAMD, CUDA, and Clusters: Taking GPU Molecular Dynamics Beyond the Desktop

A supercomputer is only as fast as its weakest link. The highly parallel molecular dynamics code NAMD was one of the first codes to run on a GPU cluster when G80 and CUDA were introduced in 2007. Now, after three short years, the Fermi architecture opens the possibility of new algorithms, simpler code, and easier optimization. Come learn the opportunities and pitfalls of taking GPU computing to the petascale.

Speaker: James Phillips, University of Illinois

Topics: Molecular Dynamics, High Performance Computing, Life Sciences, Physics Simulation

Time: Thursday, September, 23rd, 14:00 - 14:50

2088 - Nucleotide String Matching Using CUDA-Accelerated Agrep

Dive deep into the intelligent utilization of various CUDA memory spaces to remarkably speedup approximate DNA/RNA nucleotide sequence matching algorithm in bioinformatics by an amazing factor of 67 compared to multi-threaded quad core CPU counterpart. Our talk provides a very good example to demonstrate how to use indexable array to save frequently updated variables directly into GPU registers, how to organize shared memory into a 2D array to avoid bank conflict, and how to shuffle the data structure to satisfy the requirement for coalesced global memory access. Our CUDA implementation employs online approach and can be applied in real time.

Speaker: Hongjian Li, The Chinese University of Hong Kong

Topics: Life Sciences, Algorithms & Numerical Techniques

Time: Thursday, September, 23rd, 16:00 - 16:50

2248 - Parallel Processing on GPUs at the University of Utah

The University of Utah is a CUDA Center of Excellence. We have been doing both basic and applied research using CUDA. In this session, we plan to give 3-4 talks on ongoing research. Most of the work that we will be presenting has been peered reviewed at top conferences.

Speakers: Claudio Silva, University of Utah, Huy Vo, University of Utah

Topics: High Performance Computing, Life Sciences, Medical Imaging & Visualization, Tools & Libraries

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

2218 - Redesigning Molecular Dynamics for GPUs and GPU Clusters

Generalized Born and Particle Mesh Ewald (PME) molecular dynamics are two computationally intensive algorithms for simulating biological molecules. While several adaptations of Generalized Born have attained excellent speedup on GPUs, high performance Particle Mesh Ewald has been more elusive. Here we describe in detail a recent port of PME implemented within AMBER 11 that has achieved performance on par with up to 128 nodes of a top ten supercomputer.

Speaker: Scott Le Grand, NVIDIA

Topics: Molecular Dynamics, Algorithms & Numerical Techniques, High Performance Computing, Life Sciences

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

2034 - Reformulating Algorithms for the GPU

Important applications in signal, data processing and bioinformatics that use dynamic programming are difficult to parallelize due to intrinsic data dependencies. We demonstrate a novel technique to extract parallelism out of data dependent algorithms and reformulate the same for GPUs.

This simple technique breaks the dependencies and resolves them at an optimal point later in time, thus obtaining remarkable speedup on GPUs. We present a case study from computational biology i.e., protein motif-finding. We also present how the same technique can be extended and applied to other relevant problems such as gene-prediction and phylogenetics.

Speakers: Narayan Ganesan, University of Delaware, Michela Taufer, University of Delaware

Topics: Life Sciences, Algorithms & Numerical Techniques, High Performance Computing

Time: Wednesday, September, 22nd, 11:00 - 11:20

2278 - Strategies for Code Encapsulation in GPU Implementations

Code encapsulation is a common technique used to reduce code complexity that a given programmer has to understand. It allows the use of increasingly complex systems of hardware, software, and algorithms to tackle increasingly difficult scientific problems. Unfortunately, code encapsulation is not easily attainable in current GPU environments. We will share our OpenCL development experiences for achieving partial encapsulation in GPU implementations, and discuss best practices in this area.

Speaker: Brian Cole, OpenEye Scientific Software

Topics: Programming Languages & Techniques, High Performance Computing, Life Sciences

Time: Thursday, September, 23rd, 09:00 - 9:20

2242 - Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion

Ever since the 1827 discovery of Brownian motion by observing pollen grains, quantifying motion under the microscope has led to breakthroughs in physics, biology and engineering. Here, I present methods we have developed using confocal microscopy to deduce 3D structure and dynamics from 2D image sequences. We analyze the motion of diffusing colloidal particles and swarms of bacteria free to swim in 3D, which we observe at the single-organism level. We rely heavily on GPU computing to process our large data sets, making extensive use of NPP, CuFFT and optical-flow CUDA algorithms originally developed for machine vision in automobiles.

Speaker: Peter Lu, Harvard University

Topics: Computer Vision, Imaging, Life Sciences

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

2172 - Unveiling Cellular & Molecular Events of Cardiac Arrhythmias

George Mason University is using CUDA technology to get a 20x speed-up in simulations of intracellular calcium dynamics, thought to play a major role in the generation of cardiac arrhythmias. We will discuss the novel algorithms we have developed for Markov Chain Monte Carlo Simulation and their use in investigating elementary events of calcium release in the cardiac myocyte. The resulting extremely fast simulation time has generated new insights into how defects in the control of intracellular calcium may lead to cardiac arrhythmia.

Speaker: Tuan Hoang-Trong, George Mason University

Topics: Life Sciences, Algorithms & Numerical Techniques, Physics Simulation

Time: Tuesday, September, 21st, 11:00 - 11:50

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 Insitute

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

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