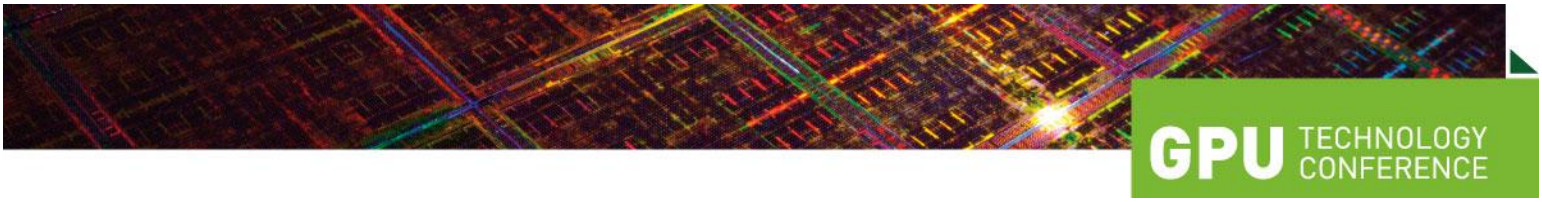

GPU Technology Conference, May 14-17, 2012
McEnery Convention Center, San Jose, California
www.gputechconf.com

Sessions on Astronomy & Astrophysics (subject to change)

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

S0618 - Best Practices of a 800TFlop Hybrid Supercomputer Implementation *(Presented by Appro)*
Steve Lyness (Appro), Taisuke Boku (University of Tsukuba)
Day: Tuesday, 05/15 | **Time:** 9:30 am - 10:20 am
Topic Areas: Supercomputing; Astronomy & Astrophysics
Session Level: Intermediate

Learn about the "Frontier Computing System", deployed by Appro for the University Of Tsukuba Center Of Computational Sciences in Japan containing over half a million GPU cores. Learn how reliability, availability, manageability and compatibility were essential for this successful 800TF hybrid supercomputing implementation. Explore new techniques in how HA-PACS is accelerating large scale parallel code by combining CPU/GPU processing cluster configurations for scientific research, such as astrophysics and climate modeling. Learn how to improve data I/O performance and memory size limitations in hybrid systems configured with Lustre™ File System offering the best performance per dollar and excellent memory capacity per/FLOP.

S0347 - Accelerating Radio Astronomy Cross-Correlation beyond 1 Tflops using Fermi
Michael Clark (NVIDIA)
Day: Thursday, 05/17 | **Time:** 9:00 am - 9:50 am
Topic Areas: Astronomy & Astrophysics; Supercomputing
Session Level: Intermediate

Radio astronomy is a signal processing application that requires extreme supercomputing. While today's radio telescopes require 10-100 Tflops of computational power, by the end of the decade this will increase to 1 Exaflops. The most compute intensive part of this problem is the so-called cross-correlation algorithm, which is a linear-algebra problem. In this session we demonstrate that the Fermi architecture is ideally suited to this problem, and through exploiting the Fermi memory hierarchy it is possible to achieve close to 80% of peak performance in a real application.

S0124 - Signal Processing on GPUs for Radio Telescopes
John Romein (ASTRON)
Day: Thursday, 05/17 **Time:** 10:00 am - 10:50 am
Topic Areas: Astronomy & Astrophysics
Session Level: Intermediate

In this talk, we will present GPU implementations of four highly compute-intensive algorithms used by radio telescopes.

S0187 - GPUs for Radio Imaging
Vamsi Krishna Veligatla (University Of Groningen)
Day: Thursday, 05/17 | **Time:** 2:00 pm - 2:25 pm
Topic Areas: Astronomy & Astrophysics
Session Level: Intermediate

With the advent of a new breed of Telescopes like the Low Frequency Array (LOFAR), which rely on software processing to process large data-sets that they generate, there is a need to improve the software to run as fast as possible in order to process the large data-sets in a reasonable time. In this session we describe how we have used the computing power of GPU's to improve the performance of the standard radio imaging techniques as well as how this computational power is useful for creating a new generation of Radio Imaging Algorithms.

S0022 - Scalable Frameworks and Algorithms for Terascale Radio Astronomy Images
Christopher Fluke (Swinburne University of Technology - Centre for Astrophysics and Supercomputing)
Day: Thursday, 05/17 | **Time:** 2:30 pm - 2:55 pm
Topic Areas: Astronomy & Astrophysics; Visualization
Session Level: Intermediate

Learn how the oldest science is using the newest processors to solve a critical problem: how to accomplish traditional image analysis and visualization tasks when the images are terabytes in size? Simple, standard operations such as displaying 2-d slices, evaluating image statistics, and applying histogram equalization become manifestly challenging when images dramatically exceed single-node memory capacity. We will explain how our hybrid CPU-GPU cluster framework - which can volume render a 200GB image at >50fps! - will support traditional radio astronomy tasks for the colossal images that the Square Kilometre Array and its precursor, the Australian SKA Pathfinder, will generate.

S0087 - GPU Acceleration of Dense Stellar Clusters Simulation
Bharath Pattabiraman (Northwestern University), Stefan Umbreit (Northwestern University)
Day: Thursday, 05/17 | **Time:** 3:00 pm - 3:25 pm
Topic Areas: Astronomy & Astrophysics; Computational Physics; Algorithms & Numerical Techniques
Session Level: Intermediate

Computing the interactions between stars within dense stellar clusters is a problem of fundamental importance in theoretical astrophysics. This paper presents the parallelization of a Monte Carlo algorithm for simulating stellar cluster evolution using programmable Graphics Processing Units. The kernels of this algorithm exhibit high levels of data dependent decision making and unavoidable non-contiguous memory accesses. However, we adopt various parallelization strategies and utilize the high computing power of the GPU to obtain substantial near-linear speedups which cannot be easily achieved on a CPU-based system. This acceleration allows to explore physical regimes which were out of reach of current simulations.

S0111 - An Efficient CUDA Implementation of a Tree-Based N-Body Algorithm
Martin Burtscher (Texas State University)
Day: Thursday, 05/17 | **Time:** 3:30 pm - 4:20 pm
Topic Areas: Application Design & Porting Techniques; Astronomy & Astrophysics; Molecular Dynamics; Supercomputing
Session Level: Advanced

This session presents a complete CUDA implementation of the irregular Barnes-Hut n-body algorithm. This algorithm repeatedly builds and traverses unbalanced trees, making it difficult to map to GPUs. We explain in detail how our code exploits the architectural features of GPUs, including lockstep operation and thread divergence, both of which are commonly viewed as hurdles to achieving high performance, especially for

irregular codes. On a five million body simulation running on a Tesla C2050, our CUDA implementation is 30 times faster than a parallel pthreads version running on a high-end 6-core Xeon.