


Session 1: GPU: Graphics Processing Units

Miriam Leeser / Northeastern University

**HPEC Conference
15 September 2010**




MIT Lincoln Laboratory




Manycore is Here to Stay

- **Current accelerators of choice:**
 - Multicore processors (2, 4 or 8 processors on a chip)
 - Graphics processing units (GPUs)
NVIDIA, AMD
 - Cloud computing
- **The future:**
 - Manycore processors (16, 32, 64 processors on a chip)
 - More flexible GPUs
AMD Fusion, NVIDIA Fermi, Intel Knight's Corner
 - Cloud computing on a chip
Intel's Single Chip Cloud Computer



This Morning's Talks

- **Invited speakers:**
 - **Accelerating Mechanical Computer Aided Engineering (MCAE) applications with GPUs**
Dr. Robert Lucas, USC ISI
 - **Thinking outside the Tera-Scale box**
Piotr Luszczek, University of Tennessee at Knoxville
- **GPUs:**
 - **Sparse Matrix Algorithms on GPUs and their Integration into SCIRun**
Devon Yablonski, Northeastern University
 - **Benchmark Evaluation of Radar Processing Algorithms on Graphics Processor Units (GPUs)**
Scott Sawyer, Lockheed Martin
 - **Failing In Place for Low-Serviceability Infrastructure Using High-Parity GPU-Based RAID**
Anthony Skjellum , University of Alabama at Birmingham



The Future

- **Manycore computing and GPUs:**
 - From dozens to thousands of cores on a single chip
 - SIMD: Single Instruction Multiple Data
 - Vector processing
 - Network on a chip
- **A mix of styles**
 - CPU, GPU, ...
- **Programming techniques, languages, methodology**
 - ?