

Two vertical lines, one blue and one orange, are positioned on the left side of the slide.

VSCSE Summer School 2008

Accelerators for Science and Engineering
Applications: GPUs and Multi-cores


Lecture 10
Future Directions and Conclusions

Why is programming many-core processors costly today?


- Separate structure from CPU
 - Data isolation and marshalling with pressure to reduce overhead
 - GMAC tool coming
- Lack of standardized programming interface
 - Each has its own app development models and tools
 - OpenCL a good start, Gluon will raise the level much higher
- Management of memory resources
 - CUDA-lite under Gluon coming
- Multi-dimensional optimizations required for achieving performance goals
 - CUDA-tune and ADAPT model tools coming
- Programming multiple GPUs and CPU cores still tedious
 - WorkForce runtime coming
- Too many people start projects from scartch
 - CUDA Research Zone coming, new application frameworks

A different approach from the past

- Simple parallelism
 - Focus on simple forms of parallelism for programmers
 - Trade some generality and performance for productivity
- Power tools
 - Leverage and strengthen app development frameworks
 - Empower tools with specification, analysis and domain information



Reduced Tuning Efforts



To Learn More

- A set of more comprehensive lectures
 - <http://courses.ece.illinois.edu/ece498/al/Syllabus.html>
(slides, voice recordings)
- A new book
 - <http://courses.ece.illinois.edu/ece498/al/Syllabus.html>
(chapters 7-9 brand new!)
- CUDA Programming Guide and Related Documents
 - http://developer.download.nvidia.com/compute/cuda/2_0/docs/NVIDIA_CUDA_Programming_Guide_2.0.pdf

A Great Opportunity for Many

- GPU parallel computing allows
 - Drastic reduction in “time to discovery”
 - 1st principle-based simulation at meaningful scale
 - New, 3rd paradigm for research: computational experimentation
- The “democratization” of power to discover
 - \$2,000/Teraflop SPFP in personal computers today
 - \$5,000,000/Petaflops DPFP in clusters in two years
 - HW cost will no longer be the main barrier for big science
 - You will make the difference!
- Join the international CUDA research, education and development community