

Unleashing the Power of the Cell BE for HPC Applications

David Bader
Georgia Tech.
bader@cc.gatech.edu

Michael Perrone
IBM Research
mpp@us.ibm.com

Ashok Srinivasan
Florida State University
asriniva@cs.fsu.edu

Panelists

- Virat Agarwal, Georgia Tech./IBM
- Scott Baden, UCSD
- David Bader, Georgia Tech.
- Alex Chow, IBM
- Michael McCool, RapidMind
- Ashok Srinivasan, Florida State

- Slides will be available from a link at sti.cc.gatech.edu

Unleashing the Power of the Cell BE for HPC

Applications: Topics

- IBM's roadmap for the Cell
- Which types of applications suit the Cell and which types suit GPUs?
- What will be required to enable the Cell to have a role in exascale computing?
- How can Cell blades be integrated with traditional Intel or AMD processor nodes in a cluster in terms of resource management (schedule) and user running environment?

Unleashing the Power of the Cell BE for HPC Applications: Topics

- What has hindered people from porting their applications to the Cell?
- What changes or additional features will induce people to use Cell for their applications?
- What are some ways in which latency and bandwidth problems can be handled?

Unleashing the Power of the Cell BE for HPC

Applications: Topics

- What is the current status of the Cell port of various standard libraries like FFT, BLAS, LAPACK, etc.
- What are the plans for stronger Fortran support on Cell?
- What development tools are available, and what are the planned improvements are planned (for example, Cell SDK, Octopiler, IBM's XL compilers, Cell SuperScalar, Sony MARS, RapidMind, etc.)?
- Which programming models for the Cell are more successful and likely to survive
- How can the Cell be made more competitive in terms of code porting efforts and cost against Intel's?