Fall Creek Falls – Session IIIa

Key Challenges presented by next generation hardware systems

Future Technologies Group ornl

Managed by UT-Battelle
for the Department of Energy

Oak Ridge National Laboratory

Session Agenda

- **Doug Burger - Texas**
  - Termites, Chainsaws, and Bulldozers: The Case for Composable Processors
- **Dimitris Nikolopoulos - Virginia Tech**
  - System Software Challenges and Opportunities on Asymmetric Multi-core Systems
- **Volodymyr Kindratenko - NCSA/UIUC**
  - High-Performance Computing on FPGAs: Challenges and Opportunities
- **John Levesque - Cray**
  - Scaling Beyond Commodity
- **Sid Chatterjee - IBM**
  - Breaking Down Walls
- **Panel**
  - All speakers

Heterogeneous Architectures

Coming Soon to a System Near You

Jeffrey Vetter

ORNL/Georgia Tech

Managed by UT-Battelle
for the Department of Energy

Best of times, worst of times…

- Years of Prosperity
  - Increasing large-scale parallelism
  - Increasing number of transistors
  - Increasing clock speed
  - Stable programming models and languages

- 'New' Constraints on architectures
  - Power
  - Heat / thermal envelope
  - Signaling
  - Packaging
  - Instruction level parallelism
  - Memory, I/O, interconnect latency and bandwidth
  - Market trends favor 'good enough' computing – *Economist*

Architectures – Multicore

- **Core counts (peak flops) are easy to increase**

1024…

160…

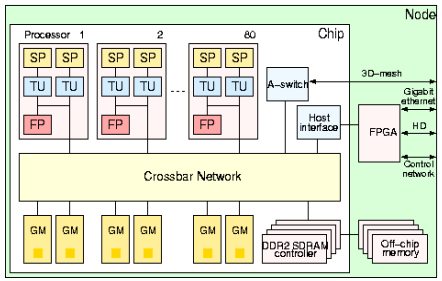

80…

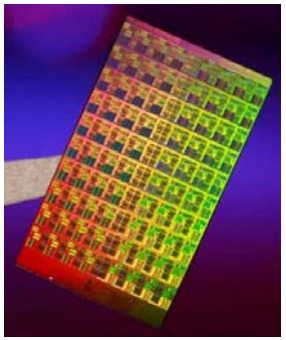

4…

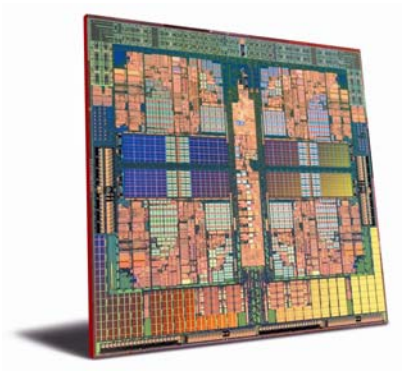

Other resources are the challenge!

Architectures – Heterogeneity, Specialization

Architectures target specific workloads: games, graphics, business, encryption, media, etc.

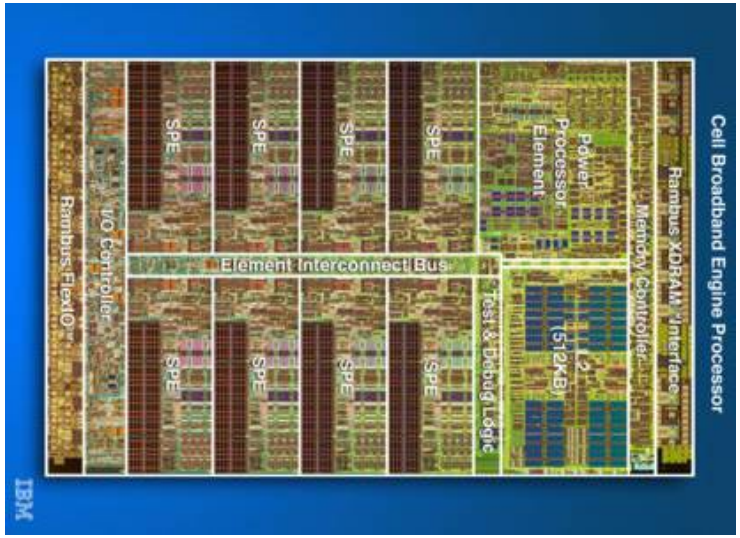

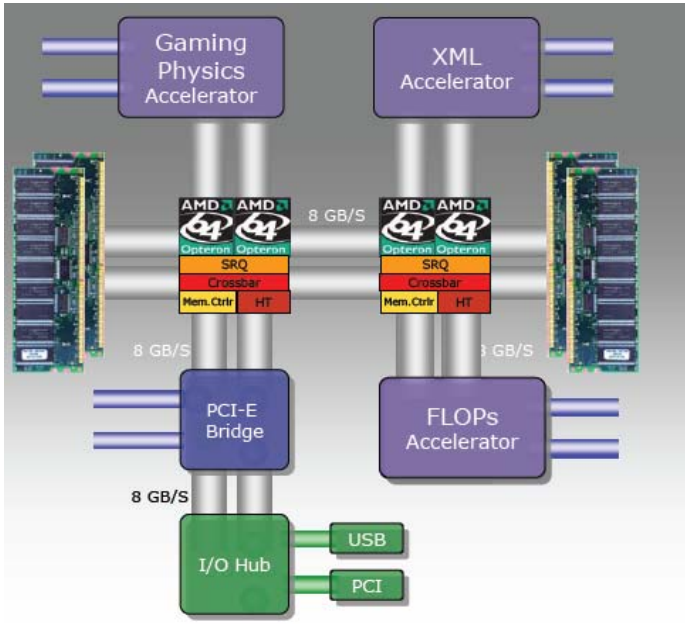

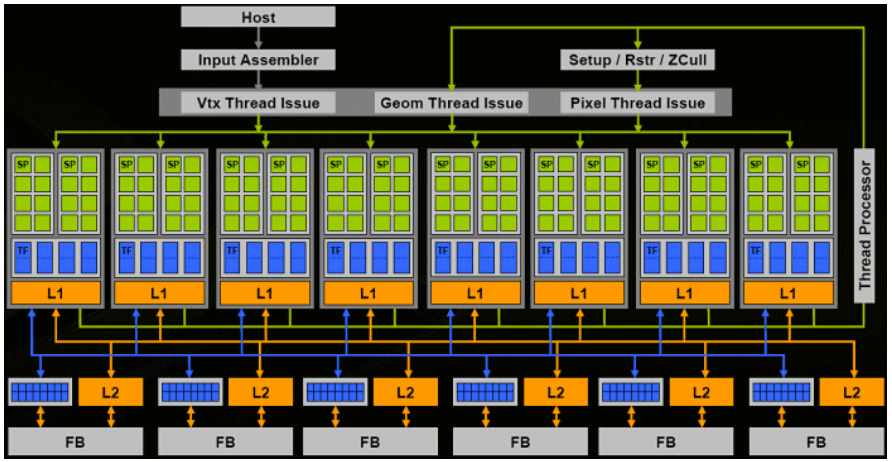

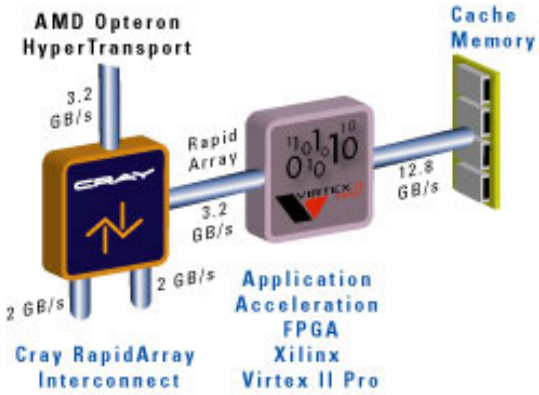

Case Study of Molecular Dynamics across Various Architectures

- Bio-molecular simulation techniques: Molecular Dynamics
- MD simulations
  - N^2 complexity (reduced by cutoff)
  - N order of 10K to over 1000K atoms
  - Unit of simulation time: femto seconds
  - Desirable simulation: with 100K atoms over 10^{-3} and 10^{-6} seconds
- Current capabilities
  - A few hundred pico-seconds per simulation day on a multi-core desktop system
  - With AMBER on MPP systems: few hundred pico-seconds on up to 128 processors
  - With LAMMPS on MPP systems: couple of nano-seconds on ~10K processors
- Focused on accelerating non-bonded calculations

J.S. Meredith, S.R. Alam, and J.S. Vetter, "Analysis of a Computational Biology Simulation Technique on Emerging Processing Architectures," Proc. HiCOMB 2007, Sixth IEEE International Workshop on High Performance Computational Biology, 2007.

Contemporary Heterogeneous Architectures

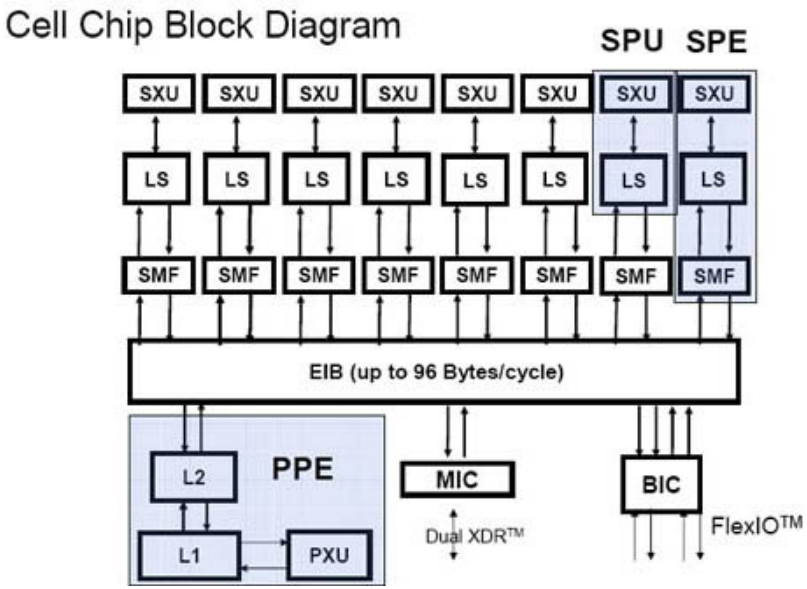

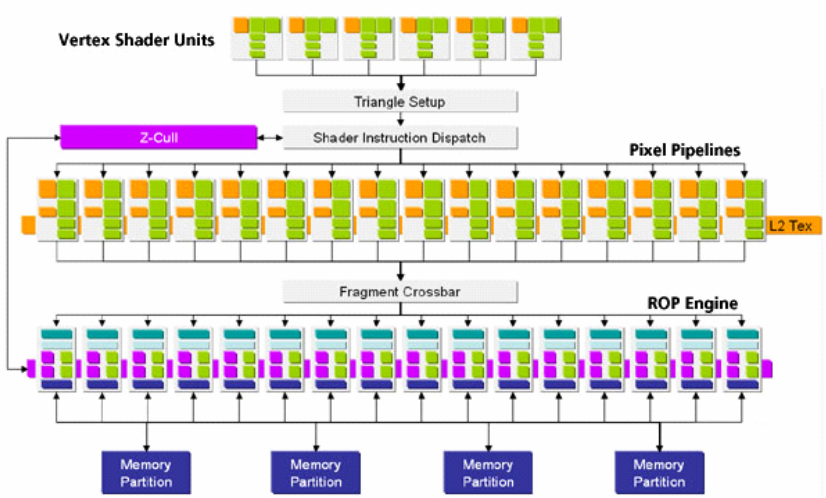

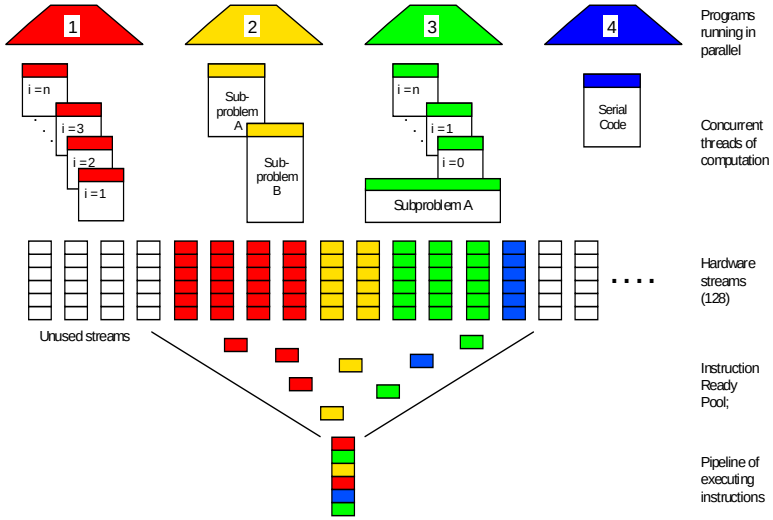

Cell

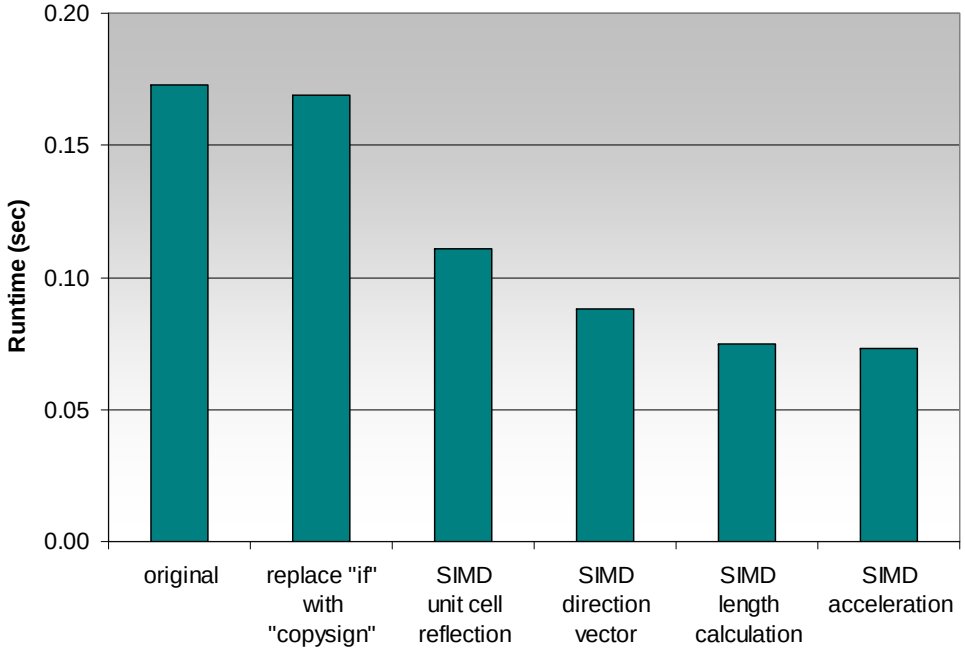

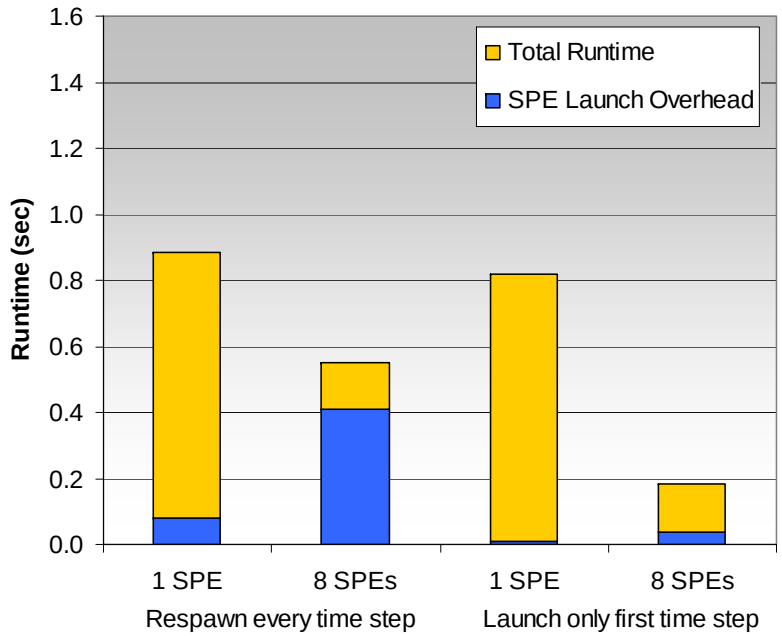

Number of Atoms	2048
Opteron	**0.925 sec**
Cell, 1 SPE	**0.816 sec**
Cell, 8 SPEs	**0.181 sec**
Cell, PPE only	**4.701 sec**

GPU

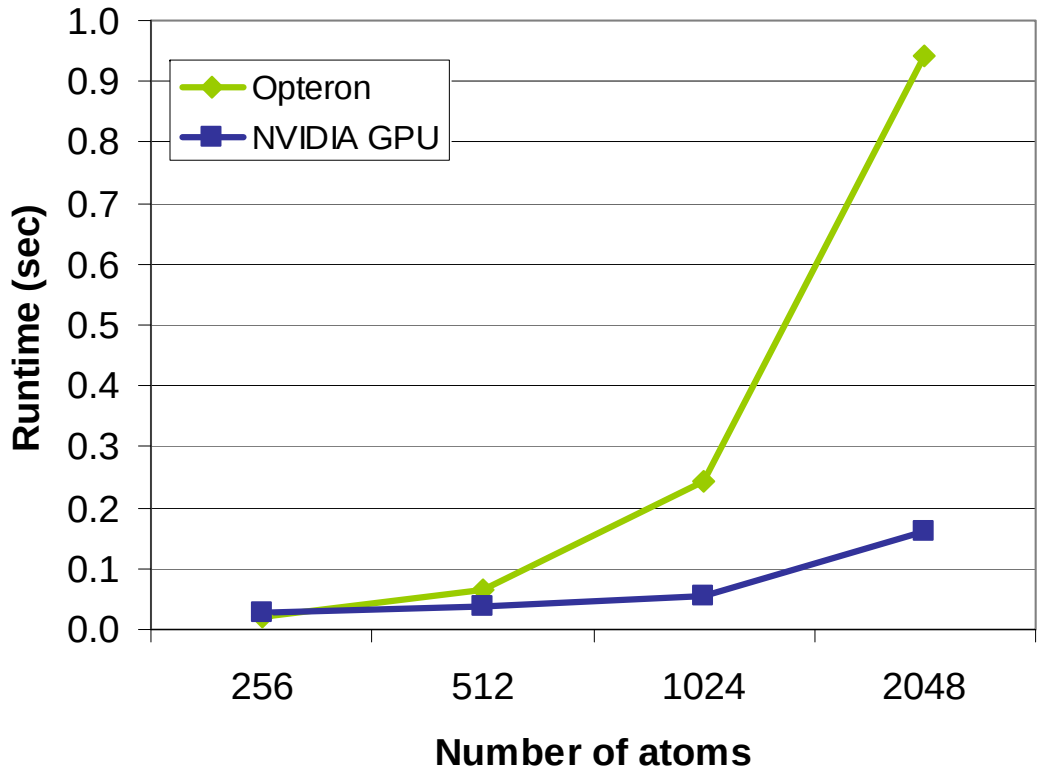

- **NVIDIA GeForce 7900GTX GPU versus a 2.2GHz Opteron**
- **6x speedup including data transfer overheads**

MTA-2

- Realities
  - 200 MHz vs. 2.2 GHz
- Opteron ~2x faster
- Compared scaling with increasing number of atoms
- Opteron runtime increases faster rate

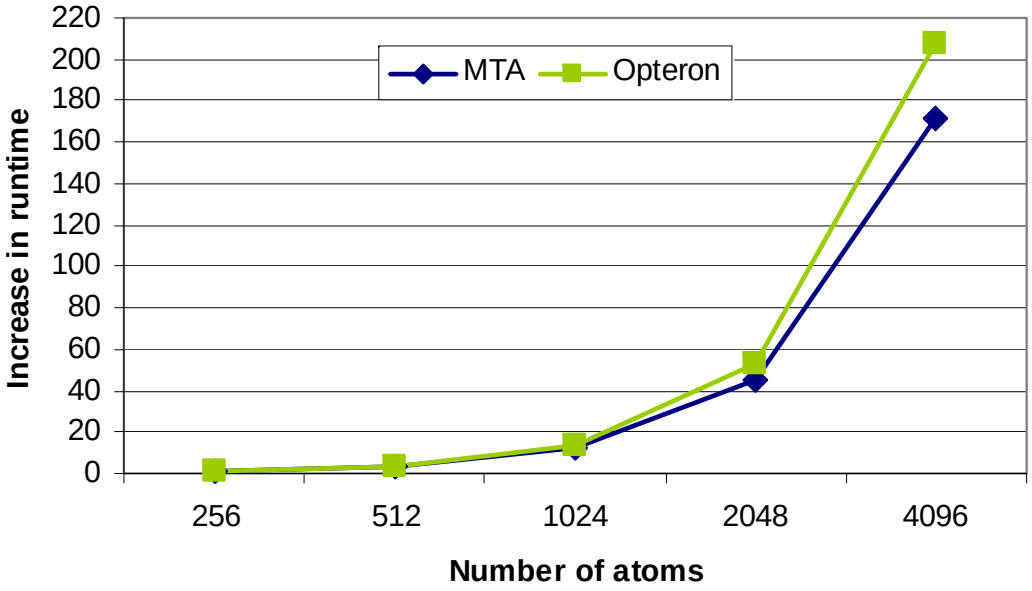

MD Conclusions

- Comparison:
  - Productivity vs. performance
  - Cost effectiveness and availability
- Source line of code comparison (SLOC)
  - GPU highest
  - MTA-II lowest
- Performance comparison (2048 atoms):
  - 6x on GPU, increase with number of atoms
  - 5x on Cell
  - 0.5x on MTA-II
- Can we sustain these for:
  - Full code acceleration
  - Double-precision calculations
- Is performance gap going to increase/decrease with multi-core systems?

Key Experiences (and Challenges)

- No common software stack across platforms
  - OS, runtime, virtualization different for each system
  - Lhotse software stack
- Performance projections
  - Uncertainty in architecture/application mapping
  - Modeling Assertions for symbolic application model
- Performance instabilities
  - Small changes in source code can cause big changes in performance
  - Improved development environments, compilers, and libraries

More details …

[1] S.R. Alam, P.K. Agarwal et al., "Achieving Supercomputer Performance for Biomolecular Simulations on Reconfigurable Systems," *IEEE Computer*, 40(2):66-73, 2007.

[2] S.R. Alam, N. Bhatia, and J.S. Vetter, "An Exploration of Performance Attributes for Symbolic Modeling of Emerging Processing Devices," Proc. 3rd International High Performance Computation Conference (HPCC), 2007.

[3] S.R. Alam, J.S. Meredith, and J.S. Vetter, "Balancing Productivity and Performance on the Cell Broadband Engine," Proc. IEEE Annual International Conference on Cluster Computing, 2007.

[4] J.S. Meredith, S.R. Alam, and J.S. Vetter, "Analysis of a Computational Biology Simulation Technique on Emerging Processing Architectures," Proc. HiCOMB 2007, Sixth IEEE International Workshop on High Performance Computational Biology, 2007.

[5] P.C. Roth and J.S. Vetter, "Intel Woodcrest: An Evaluation for Scientific Computing," Proc. 8th LCI International Conference on High-Performance Clustered Computing, 2007.

http://ft.ornl.gov/pubs

Acknowledgements

- **Thanks to contributors, sponsors**
  - **ORNL Future Tech team http://ft.ornl.gov**
  - **Experimental Computing Lab @ ORNL**
- **More information**
  - **Future Tech: Publications: http://ft.ornl.gov/pubs**
- **Acknowledgements**
  - **This research was sponsored by the Office of Mathematical, Information, and Computational Sciences, Office of Science, U.S. Department of Energy under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.**