SAND2012-2605P

Welcome to SOS16

Fess Parker Double Tree Resort
Santa Barbara, CA

Jim Ang, Ph.D.
Scalable Computer Architectures Department
Sandia National Laboratories
Albuquerque, NM

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Acknowledgements

- **Organizing Committee:**
  - Al Geist, Barney Maccabe
  - Thomas Schultess
  - Jim Ang, Ron Brightwell
- **Logistics:**
  - Dee Cadena
- **Host Organization:**
  - The Association for High Speed Computing
- **Our Sponsors**

ASC NNSA National Nuclear Security Administration

Sandia National Laboratories

SOS16 Sponsors

- Advanced Micro Devices Inc.
- Appro International
- Cray Inc.
- IBM Corp.
- Intel Corp.
- NVIDIA Corp.
- The Portland Group Inc.
- Rogue Wave Software Inc.
- Silicon Graphics International

Last Year's Invitation

Quantifying the Cost of Exascale

- Energy to Solution
- Time to Solution
- Bytes vs. FLOPS
- R&D Investments
- Technology Development Costs
- Programming Models
- Application Development Costs

March 12-15, 2012

- **Welcome Reception: Monday March 12, evening**
- **Core Workshop: Tuesday March 13 - Thursday March 15, adjourn after lunch**

ASC NNSA Sandia National Laboratories

HPC Paradigm:
Custom versus Commodity

- The last paradigm shift in HPC was the move from Cray vector supercomputers to massively parallel processor (MPP) supercomputers
- This revolutionary change was known as:
  The Attack of the Killer Micros - Eugene Brooks, LLNL
  - Founded on a philosophy of leveraging the rapid advances that were available from commodity microprocessors that rode the wave of both Moore's Law and Dennard Scaling
  - MPPs based on commodity microprocessors killed Cray Research's custom vector supercomputer business
- The ASCI Program established critical mass for this paradigm shift by investing heavily and equally in:
  - MPP application development
  - Computer science and enabling technologies
  - Large scale platforms

The HPC Paradigm is Primed for the next Major Change

- *Five + years ago commodity microprocessors began to change*
  - *Dual core processors appeared due to power and cooling limits, and commodity processors began to fall short of performance needs for HPC*
  - *Multi/many core exacerbate the memory wall/data movement problem*
  - *The result is a growing performance gap between theoretical and realized performance for our real applications*
- *Co-design is an implicit statement that multi-core commodity processors need to be redesigned with an eye towards the needs of HPC*
- *The assumption is improvements in these new processor designs will still be mainstream so the HPC community can benefit from re-designed commodity processor volumes*

The Issue / Our Challenge:

Commodity adoption of capabilities for HPC

- How are HPC co-design innovations integrated into commodity processor designs?
- The MPP HPC paradigm, while based on commodity processor designs, has minor influence on those designs
- HPC may now have an opportunity because Industry has more transistors than they know what to do with
  - Stamping out more cores that will be even more starved for data is an indication that Industry may be receptive to good ideas
  - Co-design
  - PathForward

Panel Kickoff Statements/Questions:
Different Dimensions of Exascale Cost

- Co-design is an optimization exercise
- The traditional objective function for HPC: Minimize Time to Solution
- The new objective function for Exascale HPC may be: Minimize Energy to Solution
- Objective functions versus Constraints: What are candidate options?
- Cost dimensions:
  - HW design & development - Commodity versus Custom
  - SW design & development - System SW versus Algorithms/Solvers, versus Applications
  - Application development - Porting versus starting over

SOS16 Participants

SOS16 Sponsors

We would like to express our gratitude to these companies for their generous support.

ASC NNSA Sandia National Laboratories