Exascale Computing

IESP Meeting
October 18, 2010

THE SUPERCOMPUTER COMPANY

Agenda

- Key Challenges for Achieving Exascale
- Overview of Activities and Investments in Exascale
- Thoughts on Co-Design
- Important Issues for working with Open Source Community

Key Challenges to Get to an Exascale

CRAY THE SUPERCOMPUTER COMPANY

Power

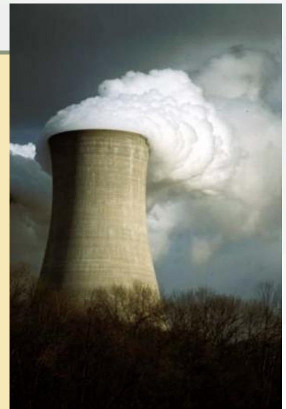

- Traditional voltage scaling is over
- Power now a major design constraint
- Cost of ownership
- Driving significant changes in architecture

Concurrency

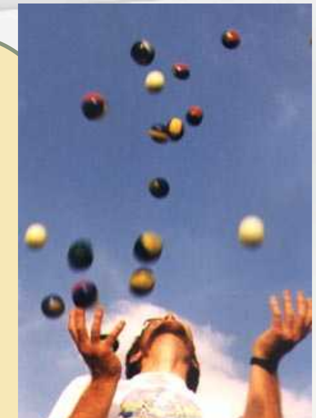

- A billion operations per clock
- Billions of refs in flight at all times
- Will require *huge* problems
- Need to exploit *all* available parallelism

Programming Difficulty

- Concurrency and new micro-architectures will significantly complicate software
- Need to hide this complexity from the users

Resiliency

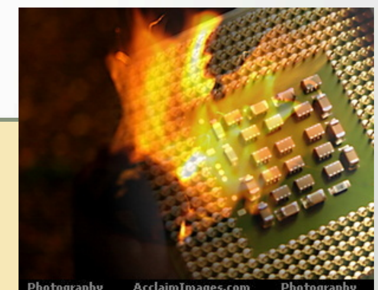

- Many more components
- Components getting less reliable
- Checkpoint bandwidth not scaling

Cray Focus Areas

- Power
  - System infrastructure, power and cooling
  - Energy-efficient network design
  - Adopting heterogeneous processor approach and actively consulting with processor vendors
- System software
  - OS and file-system scaling and jitter reduction
- Programming systems
  - Adaptive libraries and auto-tuning
  - Sophisticated runtimes for managing parallelism and locality
  - Compilers for heterogeneous processors
  - Programming tools for scoping, porting, perf analysis, and debugging
  - Languages and Programming Environments (native support for PGAS; Chapel, OpenCL, OpenMP)
- Resiliency
  - System resiliency (hardware and software support)
  - Application resiliency (strong community component)

Cray Exascale Research Initiative in Europe

Cray Media:
Nick Davis
206/701-2123
pr@cray.com

Cray Investors:
Paul Hiemstra
206/701-2044
ir@cray.com

CRAY LAUNCHES EXASCALE RESEARCH INITIATIVE IN EUROPE

Seattle, WA and Frankfurt, Germany – December 2, 2009 – As part of a company-wide goal of reaching sustained exascale performance by the end of the next decade, Global supercomputer leader Cray Inc. (Nasdaq GM: CRAY) today announced the launch of its Exascale Research Initiative at the Cray Executive Forum Europe currently taking place in Frankfurt, Germany. This research initiative will explore new ideas and technologies for overcoming the challenges of delivering a supercomputing system capable of sustained exaflop (one quintillion mathematical calculations per second) performance.

As the first company to design and build a supercomputer that achieved sustained application performance of more than one petaflops (quadrillion mathematical calculations per second), Cray is committed to the research and development of new technologies necessary to achieve exaflops computing. The challenges are significant and will require research and development into power, system and application resiliency, fault tolerant algorithms, lightweight communication and new architecture and programming models.

"We are very excited to be partnering closely with the European HPC community in kicking off this important initiative for the company," said Peter Ungaro, president and CEO of Cray. "Reaching and surpassing the petaflops barrier was an extraordinary achievement and these systems are providing unsurpassed supercomputing resources for meeting significant scientific challenges. We know there are scientific breakthroughs in important areas such as new energy sources and global climate change that are waiting for exascale performance, and we are working hard on building next-generation supercomputers that will be capable of it."

Initially, Cray will assemble a team of researchers located at the Edinburgh Parallel Computing

Cray Exascale Research Initiative in Europe – Summary

- Launched December 2009
  - *To explore new ideas and technologies for overcoming the challenges of delivering a supercomputer system capable of sustained exaflop application performance*
- Goal is to promote interaction with Application Developers
- Two collaborations with European customer sites:
  - EPCC Exascale Technology Centre, Edinburgh, UK
  - CSCS, Lugano, Switzerland

EPCC and Cray Exascale Research Initiative

- EPCC Exascale Technology Centre, Edinburgh, UK

- Projects
  - Programming models for PGAS languages
  - Programming GPU accelerators
  - Improved algorithms for FFTs
  - Network and I/O profiling

- Close Collaboration
  - Regular team meetings and calls, wiki platform

- Some Highlights
  - SC10 tutorial on Coarray Fortran
  - SC10 GPU Poster
  - OpenMP committee representation

CSCS and Cray Exascale Research Initiative

CRAY THE SUPERCOMPUTER COMPANY

- CSCS, Lugano, Switzerland
- Projects
  - Application tuning framework based on auto-tuning technology
  - GPU to GPU programming and inter-process communication
  - Programming models, tools and compilers for PGAS languages
- Close Collaboration
  - Regular team meetings and calls
- Focus on applications developed within the HP2C initiative http://www.hp2c.ch

CSCS
Swiss National Supercomputing Centre

Co-Design of Exascale Systems

- Application focus
  - Ultimate utility of the machine will be how well it runs the applications
  - Meet early and frequently with application teams to:
    - understand their scaling behavior and requirements
    - discuss potential programming models and machine features
- Labs, Cray and vendor partners define a "System Requirements" document
  - Lays out system hardware, software, and programming model characteristics
- Co-develop system concept(s) that meet the system requirements
  - Include all aspects of system design (hardware architecture, system software, programming models, resiliency, etc.)
  - Regular technical meetings to discuss alternatives
  - Continued engagements with application teams
- Refine design and create implementation specification

Important Issues for working with Open Source Community

- Software Life Cycle Management
- Test Platform and Test Plan Coordination
- Vendor Differentiation Enablement/IP Protection

Software Life Cycle Management

- Software stack requirements need to be documented
  - Functionality and performance requirements
  - Stability and scaling requirements
- Design and Plan Documents for all components
  - Clearly identify requirements being met (or not met)
  - Tightly managed (and realistic) software release schedules
- Regular communication between vendors and code teams
  - Common development platforms/tools, e.g. code repositories, bugzilla are helpful
  - Probably need a tech transfer "hub"
    - Code teams communicate in to the hub
    - Hub communicates out to the vendors
- Support infrastructure needs to be identified
  - HW/People Resources
  - Well-defined service strategy

Test Platform and Test Plan Coordination

- Common test platform (HW and SW) for all components needed

- High-Level Test Plan that identifies:
  - What is required of the component teams
  - What is required of an integrated test team
    - Component interaction testing
    - Scale testing?
  - What is expected to be the vendor's responsibility

- Test Case repositories/common tools

- Regular communication across community and vendor test teams

Vendor Differentiation Enablement/IP Protection

- Vendor Differentiation Enablement
  - APIs can allow vendors to innovate within the software stack and/or allow alternatives, e.g.
    - Compiler tools
    - Debuggers
    - Job schedulers
    - Low-level fault and RAS
  - SW components/functionality that can be disabled and/or replaced
    - Allows vendor to completely replace components
    - Side benefit: might provide a way to introduce new functionality between major releases
- Intellectual Property Protection
  - Lots of people working on the projects – can't do NDAs with everyone
  - Vendors will need to provide information on what interfaces/APIs they need without being required to provide details on what they plan to do