

Software Working Group

Chairman's Note: This document was prepared by the "software and applications" working group and was received by the entire workshop in plenary session without modification.

Findings

- What we've learned over the last 15 years
 - Successes
 - Challenges that have emerged in the last 15 years
- Current concerns and needs for the next 5-10 years
- Impact of new hardware technologies

Five improvements in the last 15 years?

- Experience with systems and some applications at large-scale
- Independent software vendors now support distributed memory programming paradigms, narrowing programming models
- Commercial interest in exploiting multicore architectures
- Opportunity to use our community experience with extreme concurrency on mesh-based applications (more generally, apps with geometric-based decompositions)
- Adoption of software components in some legacy applications has provided a route for the evolution of codes in terms of algorithms, software, and even languages

Five priority challenges – emphasizing what has changed in the last 15 years?

- Fault tolerance for hardware unreliability
- Growth and change of the memory wall problem (latency and bandwidth)
- Greater demand for application concurrency (to exist and to be exposed)
- Fragility of the software stack
- Relationship between commodity and HPC hardware and software

Five present concerns

- Independent software vendors are not thinking at the necessary scales
- Few participants from academia have access to the state of the art HPC systems
- Exhaustion of precision at new application scales
- (Validation and) Verification problem is becoming unavoidable and intractable
- Legacy and new code present some different issues (and may need different solutions)

Impact of New Hardware Technologies

- Modifications/extensions to current approaches (languages, libraries, etc.) probably sufficient
 - Software doesn't care whether digital logic uses CMOS or nanotubes
 - Completely different approaches should also be considered and supported as high risk, high payoff

Recommendations

- Reflect three major hardware directions:
 - Commodity architectures following Moore's Law to at least 2020
 - Specialty/Custom Architectures
 - Discontinuities (e.g., Quantum computer)
- 17% per year improvement (by some metric, on some classes of apps), sustained, would have great impact
 - Exponential improvement is one way to create a qualitative difference from evolutionary, quantitative improvements
 - A software roadmap could help
 - Realistic targets should be drawn from application needs
 - Metrics must be chosen carefully to ensure solutions address classes of applications, not model problems or individual apps

Software for Commodity

- Software needs to address issues of
 - Memory Wall
 - Concurrency
 - Interoperation with other tools
 - Performance comprehension
 - Faults and errors

Software for Specialty Architectures

- In addition to commodity issues,
 - Algorithms and performance comprehension for unique features in specialty architectures (e.g., techniques for streams, vectors, threads, parcels, ...)
 - System software to lower the overall cost
- Concerns about hardware
 - Scalar performance (and other Amdahl law effects)
 - More software/hardware co-design
- Adoption path for legacy software and legacy software developers
 - What are the minimal extensions to MPI, C, and Fortran to exploit unique features?
 - Analogue is vector or multithread extensions to Fortran, not a new language

Discontinuities

- Software that supports coprocessors
 - E.g., quantum co-processor attached to conventional system
 - Also applies to hybrid systems (Commodity + FPGA, PIM, ...)
- Custom languages (such as ones for quantum computers)

Working Group Members

- Bailey, David
- Bergman, Larry
- DeHon, Andre
- Foster, Michael
- Fox, Geoffrey
- Gropp, Bill
- Gustafson, John
- Hendrickson, Bruce
- Jardin, Steve
- Johnson, Fred
- Keyes, David
- McCabe, Barney
- Nichols, Jeff
- Pundit, Neil
- Saphir, Bill
- Tufo, Henry
- Williams, Colin
- Womble, David