Invited Presentation to the 2007 Workshop on the Frontiers of Extreme Computing:

Charter to the Working Groups

Dr. Thomas Sterling

Faculty, *LSU Center for Computation and Technology*
Arnaud & Edwards Professor, *LSU Department of Computer Science*
Adjunct Faculty, *LSU Department of Electrical and Computer Engineering*
Faculty Associate, *California Institute of Technology*
Distinguished Visiting Scientist, *Oak Ridge National Laboratory*

October 22, 2007

Working Groups

- Applications and Algorithms
  - End User Problems
  - Advanced algorithmic techniques
- Enabling Technologies
  - Logic
  - Memory and storage
  - Communications
- Parallel Architectures
  - System Organizations
  - Processor Core Structures
- System Software and Programming Environments
  - Operating systems
  - Runtime systems
  - Programming models, languages, and tools

Objectives of the Working Groups

- Challenges (Monday)
  - To the extreme regimes of computing
  - To the specific objectives, approaches, methods
- Opportunities (Tuesday)
  - To contribute to the realization of computing at Exascale and beyond
  - Specific possibilities: technologies, models, structures, methods
- Potential Paths Forward (Wednesday)
  - What do we need to do to realize this future
  - What are the most promising courses of action
  - How can we connect this future to our present
- Report (Thursday)

Working Group Chairs

- Applications and Algorithms
  - Chair: Horst Simon, LBL
  - Co-Chair: Mark Stalzer, Caltech
- Enabling Technologies
  - Chair: Sheila Vaydia, LLNL
  - Co-Chair: Erik DeBenedictis
- Parallel Architectures
  - Chair: Larry Bergman, JPL
  - Co-Chair: Loring Craymer, USC ISI
- System Software and Programming Environments
  - Chair: Ron Brightwell, SNL
  - Co-Chair: Jeff Vetter, ORNL

Cross Cutting Issues

- Power
  - Fundamental constraint
- Parallelism
  - Addresses increased performance and latency hiding
- Performance
  - Throughput (weak scaling) and reduced time (strict scaling)
- Execution Model
  - Paradigm for guiding and governing programming and computation
- Efficiency
  - Measure of delivered capability versus peak normalized to some measure of cost (e.g., dollars, power, floor space)
- Cost
  - Of development, deployment, programming, execution
- Reliability
  - MTBF

Strategic Issues – Applications and Algorithms WG

- Application drivers for Exascale and potential Science impact
- Application system resource requirements
- Algorithmic approaches to exploit Exascale system capability and structures

Strategic Issues – Enabling Technologies WG

- Limitations of conventional semiconductor and wired technologies and roadmaps in achieving Exascale computing – hence the need for alternatives
- Promising device technology innovations in logic, primary and secondary storage, and interconnect
- Nano-scale technology approaches and end of Moore’s Law
- Over the horizon concepts

Strategic Issues – Parallel Architectures WG

- Advanced execution models
- Many core structure and instruction sets
- Heterogeneous accelerators and controllers
- Global address space management and translation
- Locality management and latency hiding
- Synchronization and scheduling overhead
- Communication contention at all levels

Strategic Issues – System Software and Programming Environments WG

- Advanced execution models
- Parallel programming models, methods, and tools
- Resource management, allocation, and scheduling
- Mass storage and I/O management

Questions – Applications and Algorithms WG

1. What will be important end-user applications requiring Exascale computing?
2. What will be the needed memory capacity to support the range of Exascale applications?
3. How much parallelism can be exposed for these applications and is it enough to exploit billion-way parallelism architectures?
4. What are the expected I/O demands for mass storage including checkpoint/restart, initial conditions, intermediate results for runtime steering, and final results?
5. What are the requirements for and possibilities of algorithms suitable for Exascale applications?
6. For Exascale applications and their enabling algorithms what semantic constructs, basic mechanisms, and execution model attributes are needed?
7. What discoveries or advances might be possible with Exascale computing that are not possible with Petascale or below. Also, what advances will be possible by applying Exascale-capable technology to smaller systems (i.e., robots, spacecraft)?
8. What algorithm classes will become more or less important at Exascale as opposed to Petascale or lower. For example, will finite elements become more or less prevalent? How about Monte Carlo? Will larger computers be applied to larger simulations or large parameter sweeps over current-size simulations? While we're on the topic, will simulations be the prevalent application class for Exascale supercomputers?

Strategic Issues – Enabling Technologies WG

- What are achievable clock rates or switching times at reasonable power consumption of competing logic technologies?
- What are the contending technologies that will lead to nano-scale and the end of Moore’s law?
- What are the likely paths to intra-system communications including chip-to-chip and system-wide interconnections?
- If memory capacity will be one of the major constraining factors to application performance, what are the likely technologies and their operational characteristics towards Exascale and beyond?
- Is there life after disks? With the monotonically and potentially exponential growth in storable information, what is the likely or potential media technologies that will get us through the 21st century?
- What information processing structures may come to offer advantage over Boolean logic gates. Example: alternatives circuits, non charge-based information, quantum effects, ...
- Is there a future in systems that blend information processing and non-volatile memory? Examples may include artificial neural networks and molecular-scale FPGAs where the same physical structures learn at a low speed and process at a higher speed.
- How should quantum computing be integrated with the general information processing? What about quantum effects short of a traditional quantum computer?

Questions – Parallel Architectures WG

- What architecture strategies will enable Exascale system implementation?
- What architecture approaches will address challenges of latency, overhead, contention, and starvation expected at Exascale system structures?
- It is estimated that multi billion-way parallelism (> 100 million cores) will be required for Exascale systems. What architectures can be expected to deliver, exploit, and coordinate such unprecedented parallelism?
- What architecture techniques will minimize power consumptions?
- Are new execution models required to guide architecture structure and operations, and if so what?
- In the new Exascale architectures, how will memory and address spaces be handled?
- Will the future hold more highly parallel traditional architectures or something new. To be specific, the traditional architecture is a homogeneous parallel system comprising x86-compatible processors with MMU and a large cache. Something new might be a heterogeneous system of heavy and light cores, where the light cores are GPUs, cell, attached neural networks, etc.

Questions – System Software and Programming Environments WG

- What new semantic constructs and execution models will be needed by Exascale applications and algorithms? What will the programming languages of the future look like?
- How will programmers contend with billion-way parallelism? What will the new languages look like?
- What operating system organizing strategy will effectively manage 100 million or more processing cores efficiently and reliably?
- How will compilers and runtime systems support the new classes of applications dominated by dynamic meta data structures?
- What will be the balance in the future between user direct control of resources and system automation for ease of use?
- Will programmers continue to program with arithmetic statements, or may a different paradigm become prevalent? Examples, neural networks, fuzzy logic, graph algorithms, image processing, real time?

Applications and Algorithms

- David Bailey
- Rupak Biswas
- Phil Jones
- David Keyes
- David Koester
- Bob Lucas
- Nathan Price
- John Shalf
- Horst Simon
- Mark Stalzer
- Rick Stevens
- David Womble
- Jeff Nichols

Enabling Technologies

- George Adams
- Gian-Luca Bona
- George Bourianoff
- Erik DeBenedictis
- Alex Dickinson
- Patrick Geoffrey
- Todd Hylton
- Tony Lentine
- Jag Shah
- Thomas Theis
- Sheila Vaidya
- Stan Williams

Parallel Architecture

- Maciej Brodowicz
- Loring Craymer
- Sudip Dosanjh
- John Gustafson
- Gary Hughes
- Fred Johnson
- Rich Murphy
- Mike Niemier
- James Peery
- Steve Poole
- Karu Sandaralingam
- Steve Scott

System Software & Programming

- Ron Brightwell
- Almadena Chtchelkanova
- Guang Gao
- Mary Hall
- Fred Johnson
- Ron Minnich
- Jose Munoz
- David Probst
- Neil Pundit
- Steve Scott
- Thomas Sterling
- Jeff Vetter
- Kathy Yelick

Products, Output, & Report

- Products
  - Findings
  - Open Issues
  - Recommendations
- Output
  - Slides of Report
  - Notes by co-chair
- Report
  - Thursday morning
  - 20 minute briefing
  - Integrated by committee into single
  - Monograph: “Enabling Technologies for Exaflops Computing”

Working Group Presentations

- Monday
  - Applications and Algorithms
    - Bob Lucas
  - Enabling Technologies
    - TBD
  - Parallel Architectures
    - Steve Poole
  - System Software and Programming Environments
    - Ron Minnich
- Tuesday
  - Applications and Algorithms
    - David Bailey
  - Enabling Technologies
    - TBD
  - Parallel Architectures
    - John Gustafson
  - System Software and Programming Environments
    - David Probst
    - Guang Gao
- Wednesday
  - Applications and Algorithms
    - Volunteered Perspectives
  - Enabling Technologies
    - Volunteered Perspectives
  - Parallel Architectures
    - Volunteered Perspectives
  - System Software and Programming Environments
    - Volunteered Perspectives
- Thursday
  - Applications and Algorithms
    - Horst Simon
  - Enabling Technologies
    - Sheila Vaidya
  - Parallel Architectures
    - Larry Bergman
  - System Software and Programming Environments
    - Ron Brightwell

Pitfalls and Promise

- Please, please, please …
- Don’t solve everyone else’s problems
- Focus on your own working group!
- This is not a Teraflops workshop except as experience directly relates to Exascale & beyond

- These findings complement those of other activities
  - DARPA Exascale Study (P. Kogge)
  - DOE Exascale Townhall Meetings (E3GPS)
- “Enabling Technologies for Exaflops Computing”
  - Capture insights and Convey perspectives
  - Merge findings of all three forums