

Objectives and organization of the meeting – The Software Roadmap

Satoshi Matsuoka

Tokyo Institute of Technology

IESP Meeting Tsukuba, October 2009

Petascale Machines Circa 2010-13

Inst/Agency/Country	Name	Machine	Perf
ORNL/DoE/US 2009	Jaguar Upgrade	Cray XT5	~2PF
Utennessee/NSF/US	Cracken	Cray XT5	1PF
LLNL/DoE/US	Sequoia Proto	IBM BG/P	~1PF
Tokyo tech./MEXT/JP	TSUBAME2.0	GPU Cluster/TBD	3PF
LBNL/DoE/US 2010	Franklin 6	Cray XT6	1.2PF
Pittsburgh SC/NSF/US	???	SGI UV	2PF?
LANL/DoE/US	???	???	???
EU PRACE Machines	???	IBM/Cray/Sun/Bull...	1-2PF?
ORNL/DoE/US	Jaguar Upgrade	Cray XT6 +GPU?	20PF
NCSA/NSF/US	Blue Waters	IBM Power7 server	10-20PF
LLNL/DoE/US 2011-12	Sequoia	IBM BG/Q / PERCS	22PF
ArgonneNL/DoE/US	???	IBM BG/Q / PERCS	~20PF
Japanese-Riken/MEXT/JP	???	Fujitsu Venus	~10PF
EU PRACE Machines	???	IBM, Cray	~10PF x 4~5
China	5-6 Installations	???.Dawning?	~1PF x 6

ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems

Peter Kogge, Editor & Study Lead
Keren Bergman
Shekhar Borkar
Dan Campbell
William Carlson
William Dally
Monty Denneau
Paul Franzone
William Harrod
Kerry Hill
Jon Hiller
Sherman Karp
Stephen Keckler
Dean Klein
Robert Lucas
Mark Richards
Al Scarpelli
Steven Scott
Allan Snavely
Thomas Sterling
R. Stanley Williams
Katherine Yelick

September 28, 2008

This work was sponsored by DARPA IPTO in the ExaScale Computing Study with Dr. William Harrod

Exa-scale Computational Resources

(slide courtesy Martin Savage)

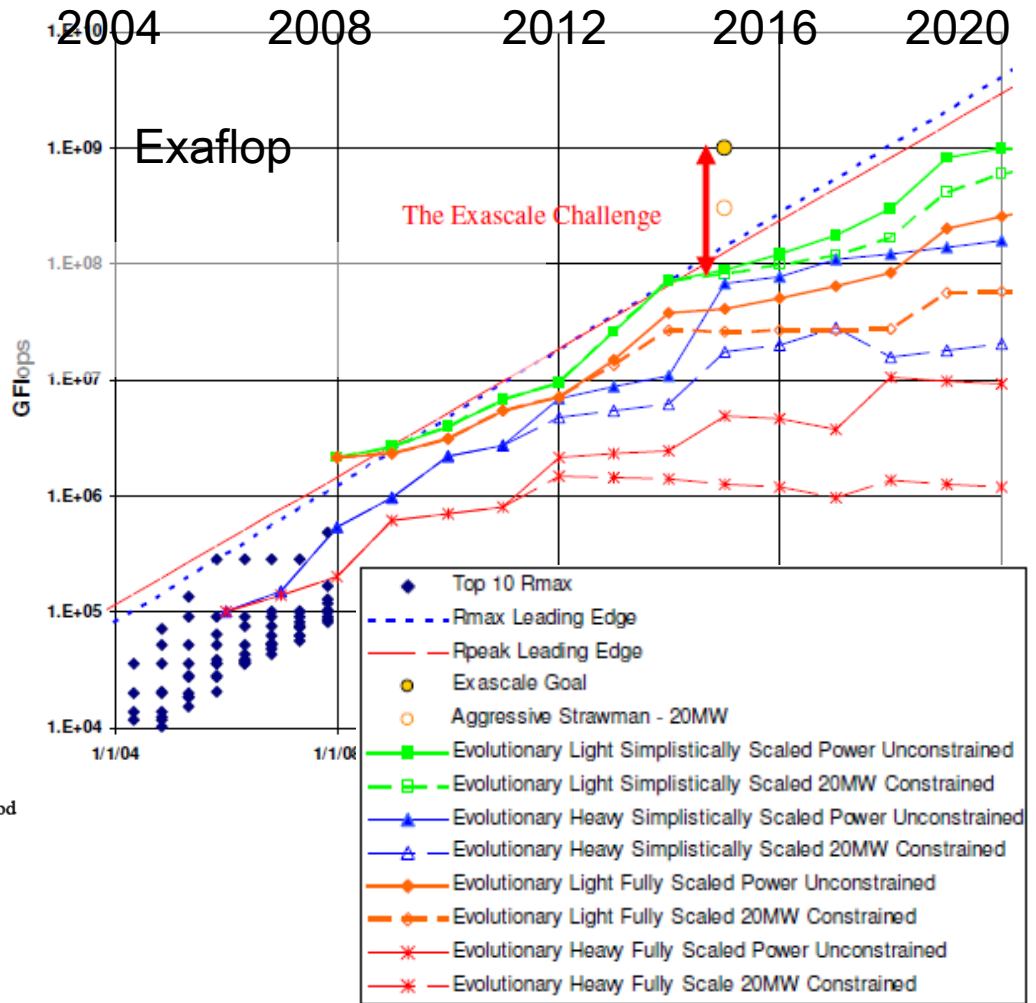
- Meeting structural and physics areas of effort

6 Application Area Exascale Workshop (2008-2009)

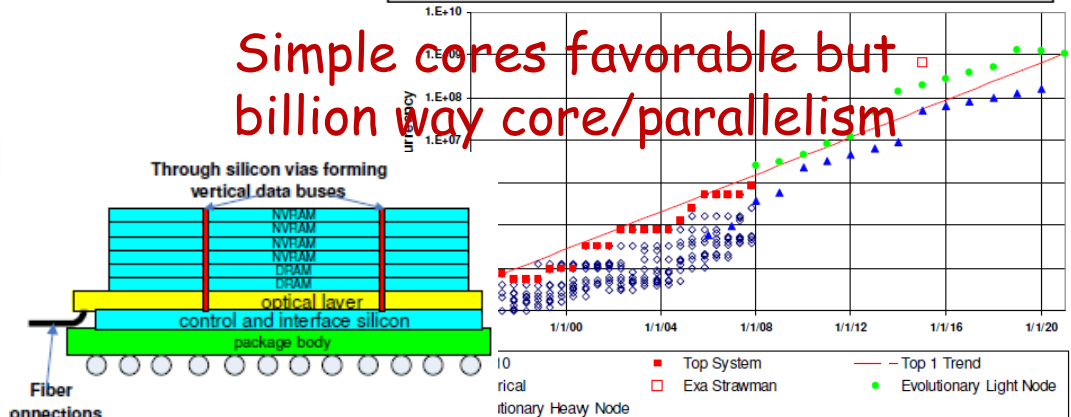


- Nuclear Astrophysics
 - Cold QCD and Nuclear Forces
 - Nuclear Structure and Reactions
 - Accelerator Physics
 - Hot and Dense QCD
- Exa-scale computing is **REQUIRED** to accomplish the Nuclear Physics mission in each area
 - Staging to Exa-flops is crucial:
 - 1 Pflop-yr to 10 Pflop-yrs to 100 Pflop-yrs to 1 Exa-flop-yr (sustained)

Peter Kogge DoD Exascale Report (SW report forthcoming)



Simple cores favorable but
billion way core/parallelism



Overall Agenda

- The 1st day there will be just two groups---SW and apps. the funding and vendor group will participate in the two groups
 - A few people will visit MEXT and Riken Peta office
- On the 2nd day the funding and vendor group will meet on their own
- They will do their report out on the 3rd day.

SW Group Agenda

- Start with review of Paris meeting
- Then overview the current draft of the roadmap document
- Next, do a brief group-level discussion of the each component
- Adjust / solicit assignments of people who have not been assigned
- Break out into small groups
- Second day intermix with application group and adjust the roadmap metrics
- Final day report out, present the necessary deliverables
- Until SC09 complete the documents

SW Group Agenda (2)

- We will spend the first day looking at both the Paris meeting and the roadmap doc
- Roadmap doc: look at the technology trends
 - We then elect leaders for each section of the roadmap
 - They will be responsible for handing back the 3 page document plus graphs/roadmap by
- Image of the deliverables
 - The “picture” with roadmap example items should be generated.
 - The docs can be in a rough form, mostly consisting of bullet items
 - Fill in the quad chart template
 - Need to be finished by Nov. SC09 BoF presentation on Wed., Nov.

Assignments of SW Components Section

- Pre-assigned two people who are the initial candidates to take ownership in finishing the document by SC09.
- Feel free to change/decline the assignments if you feel a better fit, or too busy to commit to the SC09 timeline
- We welcome additional people who would volunteer to take co-ownership

- External Env (Grid) Giovanni Aloisio
- OS Pete Beckman, Barney MacCabe
- Compilers: Barbara Chapman, Mitsuhsa Sato
- I/O: Alok Choudhary, Yutaka Ishikawa
- Libraries: Jack Dongarra, Anne Trefethen
- Algorithms: Bill Gropp, Fred Streit
- Frameworks: Mike Heroux, Robert Harrison
- Runtime: Jesus Labarta, Rajeev Thakur
- Applications Pioneers: Bill Tang, Richard Kenway
- Data & Vis: John Taylor, Rick Stevens
- System Management: Bill Kramer, Bob Wisniewski
- Debugging: Wolfgang Nagel, David Skinner
- Resilience Crosscut: Franck Cappello, Sudip Dosanjh
- Power Crosscut: Satoshi Matsuoka, John Shalf
- Performance Crosscut: Bernd Mohr, Jeffrey Vetter
- Programmability Crosscut: Thomas Sterling, Hiroshi Nakashima

Apps Group

- 1st day part one---Application group refresh themselves where they are, summary of the DoE meeting, refreshing the new participants
- 1st day part two --- walking through the tech roadmaps from the viewpoints of science and trying to come up with first cut on the scaling requirements

Apps Group Task

- Software oriented targets need to be inferred from the applications requirements
- However, given that we don't have those yet, on the first day we start from the current roadmap that has been articulated from the bottom up architectural perspectives
- On the second day the two groups gather and try to reconcile initially, and make adjustments

Apps Group Articulation Example

- Application people look at the strawman and articulate the needs without going deep into the technologies.
 - E.g. We need 10 billion grid points at so and so date so the SW needs to scale to that
 - E.g., we need shared memory programming with accelerators because that is the way our code is written.
 - Different software elements would adapt differently.
- They will also look at the roadmap document

Priority Research Direction (use one slide for each)

Key challenges

Brief overview of the barriers and gaps

Summary of research direction

What will you do to address the challenges?

Potential impact on software component

What capabilities will result?

What new methods and components will be developed?


Potential impact on usability, capability, and breadth of community

How will this impact the range of applications that may benefit from exascale systems?

What's the timescale in which that impact may be felt?

4.x <component>

<single short description of the area>



4.x <component>

- Technology drivers
- Alternative R&D strategies
- Recommended research agenda
- Crosscutting considerations