
# SC09 BOF

http://www.exascale.org

Pete Beckman & Jack Dongarra

# IESP BOF



- Overview of IESP, Jack Dongarra
- Comments from DOE, Michael Strayer
- Comments from NSF, Ed Seidel
- Comments from EC, Kostas Glinos
- Building the IESP Roadmap, Pete Beckman and Satoshi Matsuoka
- Application's Case for Exascale, Paul Messina
- European Exascale Software Initiative, Jean-Yves Berthou

# Science Case For Exascale

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

ORNL/TM-2007/238

Scientific Application Requirements for Leadership Computing at the Exascale

National Center for Computatio
December 2007

Modeling and Simulation at the Exascale for Energy and the Environment

Co-Chairs:

- DOE Exascale Steering Committee
  - ANL, LANL, LBNL, LLNL, SNL, ORNL + PNL, BNL
  - Charter: Decadal plan to provide exascale applications and technologies for DOE mission needs
- Workshops @ ~100 People
  - Climate Science (11/08)
  - High Energy Physics (12/08)
  - Nuclear physics (1/09)
  - Fusion Energy (3/09)
  - Nuclear Energy (5/09)
  - Biology (8/09)
  - Basic Energy Science (8/09)
  - Joint National Security (10/09)
  - Computer Science
  - Mathematics
  - Computer Architecture

Strong science case for the continued escalation of high-end computing.

Broad consensus necessitate the redesign and replacement of many of the algorithms and software infrastructure that HPC has built on for more than a decade.

www.exascale.org

September 28, 2008

This work was sponsored by DARPA IPTO in the ExaScale Study Program Manager; AFRL contract number FA8650-07-C-7724. of scientific and technical information exchange and its publicatio Government's approval or disapproval of its ideas or findings

NOTICE

Using Government drawings, specifications, or other data include purpose other than Government procurement does not in any way

Developing a high performance computing / numerical analysis roadmap

HPC/NA

Applications/Algorithms Roadmapping Activity

Roadmap Version 1.0

January 2009

# Looking at the Gordon Bell Prize

(Recognize outstanding achievement in high-performance computing applications and encourage development of parallel processing )



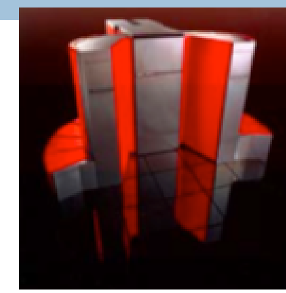

- 1 GFlop/s; 1988; Cray Y-MP; 8 Processors
  - Static finite element analysis

- 1 TFlop/s; 1998; Cray T3E; 1024 Processors
  - Modeling of metallic magnet atoms, using a variation of the locally self-consistent multiple scattering method.
- 1 PFlop/s; 2008; Cray XT5; $1.5 \times 10^5$ Processors
  - Superconductive materials

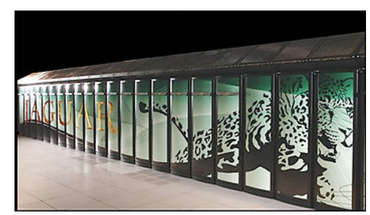

- 1 EFlop/s; ~2018; ?; $1 \times 10^7$ Processors ($10^9$ threads)

# Performance Development in Top500

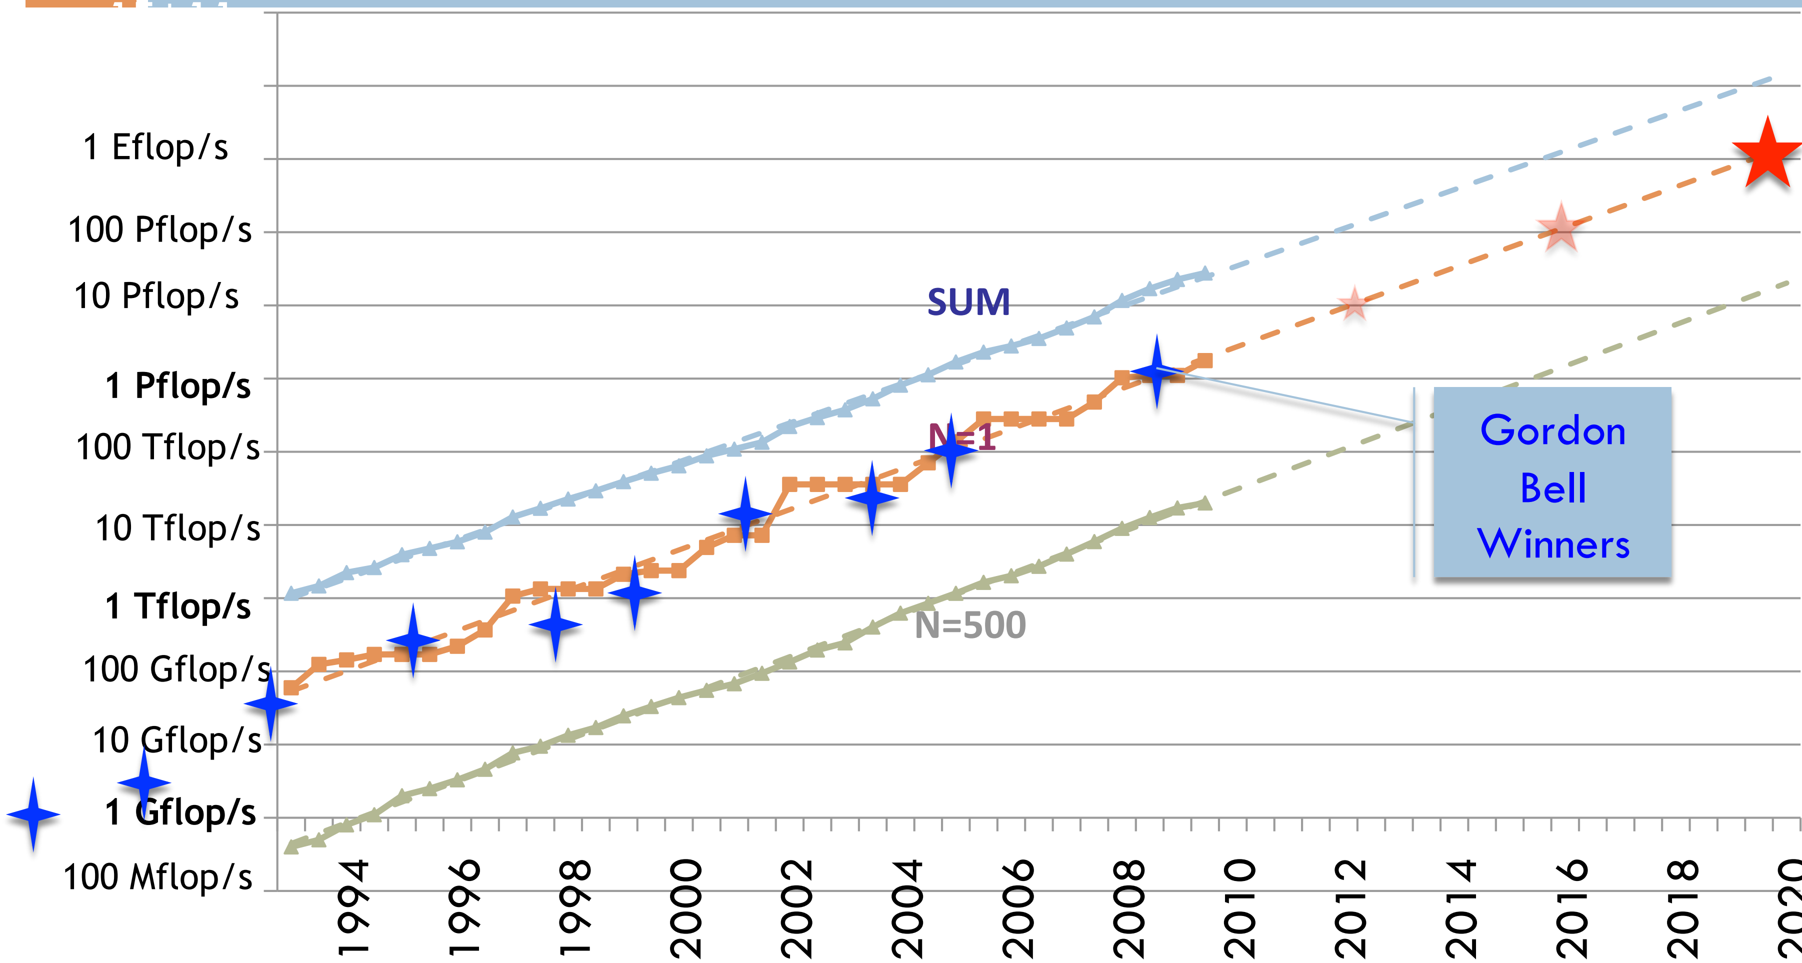


www.exascale.org

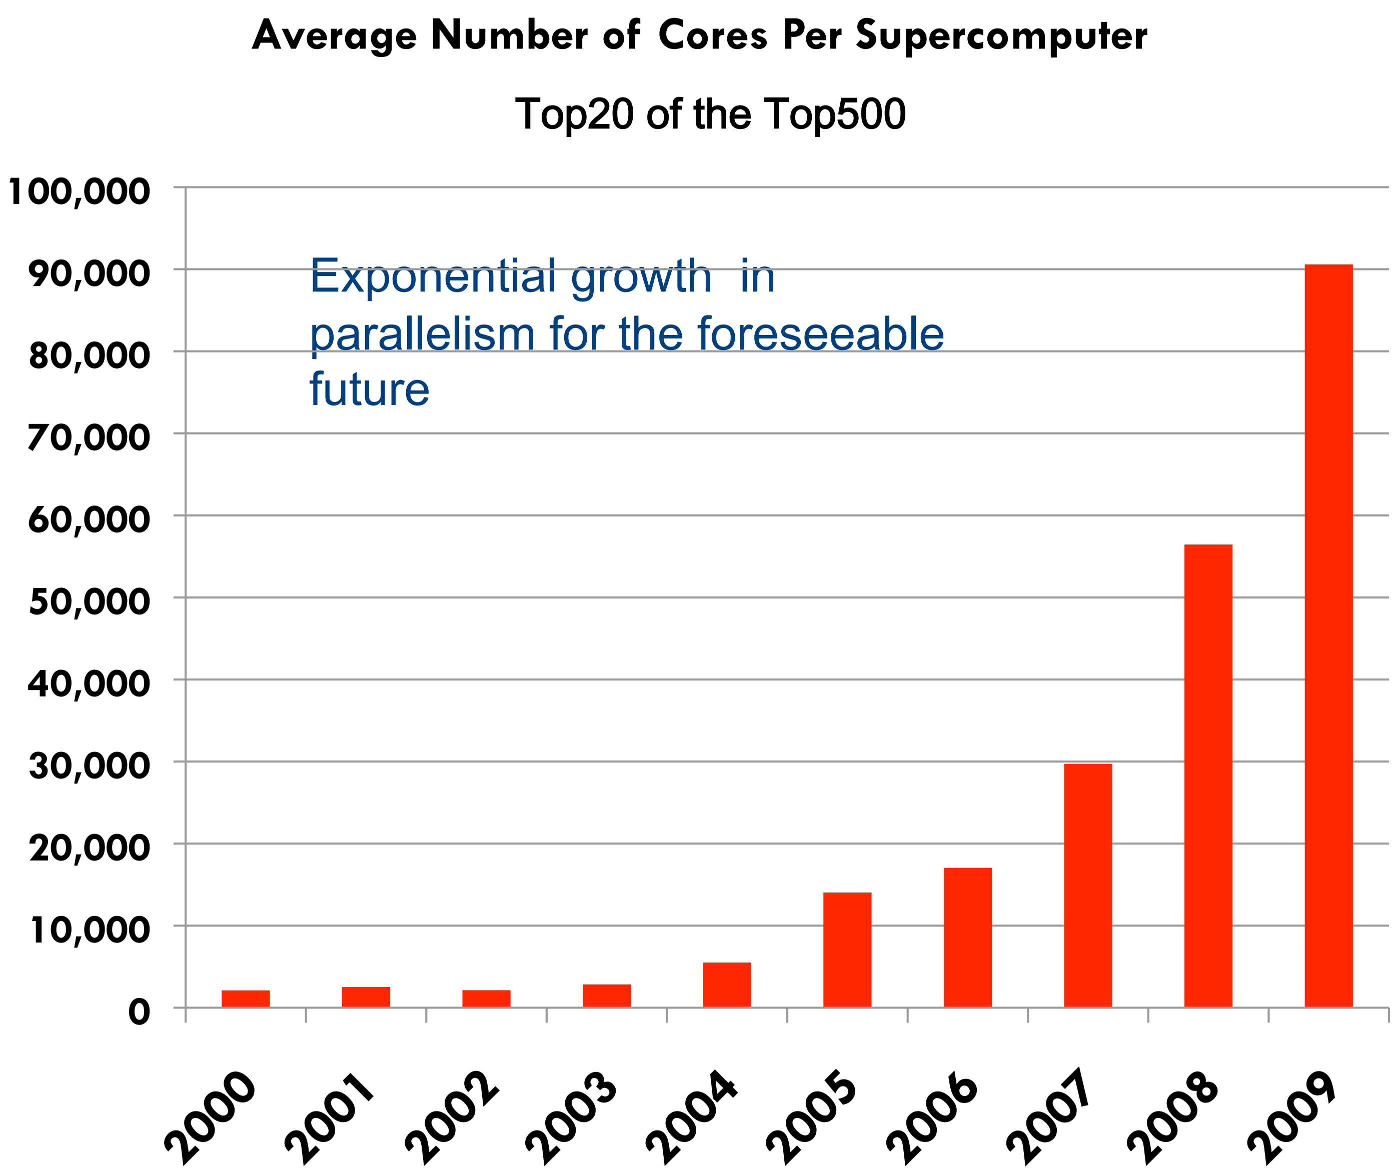
Average Number of Cores Per Supercomputer
Top20 of the Top500
100,000
90,000
80,000
70,000
60,000
50,000
40,000
30,000
20,000
10,000
0
Exponential growth in parallelism for the foreseeable future
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009

# Factors that Necessitate Redesign



- **Steepness of the ascent from terascale to petascale to exascale**
- Extreme parallelism and hybrid design
  - Preparing for million/billion way parallelism
- Tightening memory/bandwidth bottleneck
  - Limits on power/clock speed implication on multicore
  - Reducing communication will become much more intense
  - Memory per core changes, byte-to-flop ratio will change
- Necessary Fault Tolerance
  - MTTF will drop
  - Checkpoint/restart has limitations

www.exascale.org

- **Software infrastructure does not exist today**

# Exascale Computing

- Exascale systems are likely feasible by 2017±2
- 10-100 Million processing elements (cores or cores) with chips perhaps as dense as socket, clock rates will grow more slowly
- 3D packaging likely
- Large-scale optics based interconnects
- 10-100 PB of aggregate memory
- Hardware and software based fault management
- Heterogeneous cores
- Performance per watt — stretch goal 100 GF/watt of sustained performance ⇒ >> 10 – 100 MW Exascale system
- Power, area and capital costs will be significantly higher than for today's fastest systems

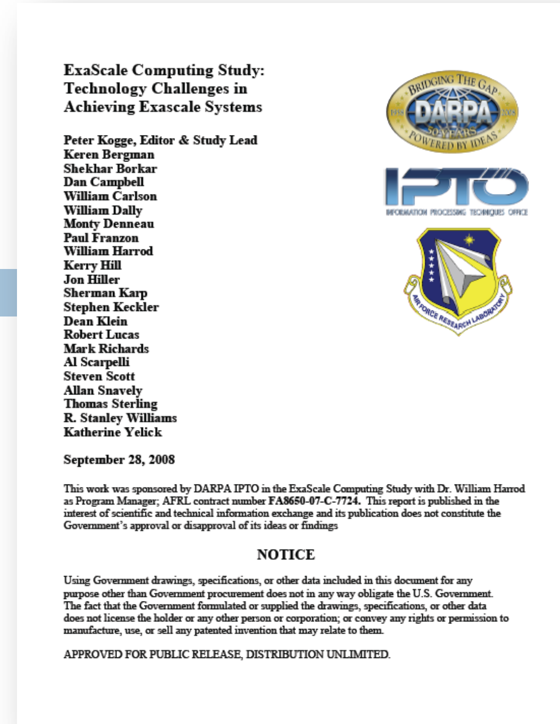

**ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems**

**Peter Kogge, Editor & Study Lead**
**Keren Bergman**
**Shekhar Borkar**
**Dan Campbell**
**William Carlson**
**William Dally**
**Monty Denneau**
**Paul Franzon**
**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 as Program Manager; AFRL contract number **FA8650-07-C-7724.** This report is published in the interest of scientific and technical information exchange and its publication does not constitute the Government's approval or disapproval of its ideas or findings

**NOTICE**

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED.

Google: exascale computing study

# A Call to Action

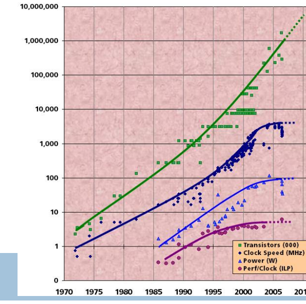




- Hardware has changed dramatically while software ecosystem has remained stagnant
- Previous approaches have not looked at co-design of multiple levels in the system software stack (OS, runtime, compiler, libraries, application frameworks)
- Need to exploit new hardware trends (e.g., manycore, heterogeneity) that cannot be handled by existing software stack, memory per socket trends
- Emerging software technologies exist, but have not been fully integrated with system software, e.g., UPC, Cilk, CUDA, HPCS
- Community codes unprepared for sea change in architectures
- No global evaluation of key missing components

# IESP Goal



Improve the world's simulation and modeling capability by improving the coordination and development of the HPC software environment

Workshops:

**Build an international plan for developing the next generation open source software for scientific high-performance computing**

# International Community Effort



- We believe this needs to be an international collaboration for various reasons including:
  - The scale of investment
  - The need for international input on requirements
  - US, Europeans, Asians, and others are working on their own software that should be part of a larger vision for HPC.
  - No global evaluation of key missing components
  - Hardware features are uncoordinated with software development

# IESP Executive Committee



- Jack Dongarra, Pete Beckman, Patrick Aerts, Frank Cappello, Thomas Lippert, Satoshi Matsuoka, Paul Messina, Anne Trefethen, Mateo Valero

# Where We Are Today:



- SC08 (Austin TX) meeting to generate interest
- Funding from DOE's Office of Science & NSF Office of Cyberinfratructure and sponsorship by Europeans and Asians
- US meeting (Santa Fe, NM) April 6-8, 2009
  - 65 people
- NSF's Office of Cyberinfrastructure funding
- European meeting (Paris, France) June 28-29, 2009
  - 70 people
  - Outline Report
- Asian meeting (Tsukuba Japan) October 18-20, 2009
  - Draft roadmap
  - Refine Report
- SC09 (Portland OR) BOF to inform others
  - Public Comment
  - Draft Report presented

**Nov 2008**

**Apr 2009**

**Jun 2009**

**Oct 2009**

**Nov 2009**

# International Journal on High Performance Computing Applications



- Set of articles that look at various aspects of Exascale software


- http://hpc.sagepub.com/current.dtl