



The Grand Challenge Question for Performance Evaluation of HPC Systems

Horst D. Simon

Director, NERSC Center and Computational
Research Divisions

Lawrence Berkeley National Laboratory

May 5, 2003

The GC Question



What quantitative and objective measure of enabling scientific discovery can be used as the basis of evaluating computing facilities for basic science?

NERSC 3 (Seaborg) Upgrade to 10 Tflop/s Completed



✍ System Characteristics:

- ✍ 416 16 way Power 3+ nodes with each CPU at 1.5 Gflop/s
 - 380 for computation
- ✍ 6,656 CPUs – 6,080 for computation
- ✍ Total Peak Performance of 10 Teraflop/s
- ✍ Total Aggregate Memory is 7.8 TB
- ✍ Total GPFS disk will be 44 TB
 - Local system disk is an additional 15 TB
- ✍ Combined SSP-2 measure is 1.238 Tflop/s
- ✍ In production now; largest unclassified system in the U.S.



Goal of NERSC: Enabling Scientific Discoveries



Borrill (LBNL) + CalTech + others.
April 27, 2000

- ✍ BOOMERANG Experiments – analyze cosmic microwave background radiation data to obtain a better understanding of the universe
- ✍ The data analysis provides strong evidence that the geometry of the universe is flat
- ✍ Computational capability provided on NERSC platforms
- ✍ MADCAP software developed at NERSC for general community

Enabling Scientific Discoveries



- ✍ Many other examples available spanning entire history of NERSC
- ✍ What are the characteristics of NERSC that enabled this work to be done here?
- ✍ What type of measures could we have used to have predicted the enablers for these scientific discoveries?

The Current Political Debate about Performance



The good news: the arrival of the Earth Simulator has made “Washington” realize that peak performance alone alone is not the answer

The bad news: “sustained to peak ratio” has replaced peak performance as the new single figure of merit

How SciDAC was “sold”



Peak Performance is skyrocketing

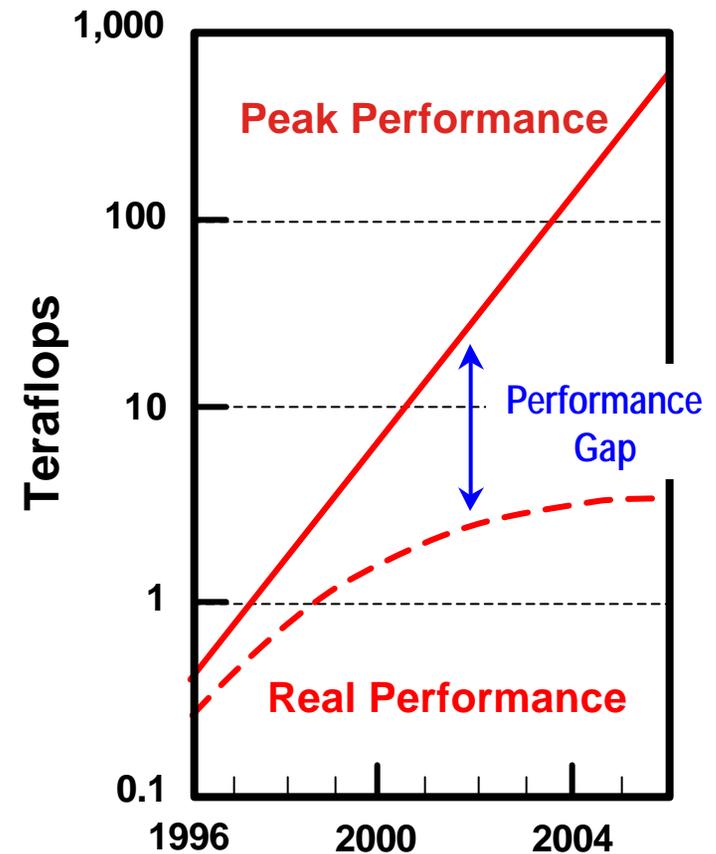
- ✍ In 1990's, peak performance increased 100x; in 2000's, it will increase 1000x

But ...

- ✍ Efficiency declined from 40-50% on the vector supercomputers of 1990s to as little as 5-10% on parallel supercomputers of today

Close the gap through ...

- ✍ Mathematical methods and algorithms that achieve high performance on a single processor and scale to thousands of processors
- ✍ More efficient programming models for massively parallel supercomputers
- ✍ Parallel Tools



Source: DOE presentation about SciDAC, 2000

... two years later ...



... you get what you set out to measure

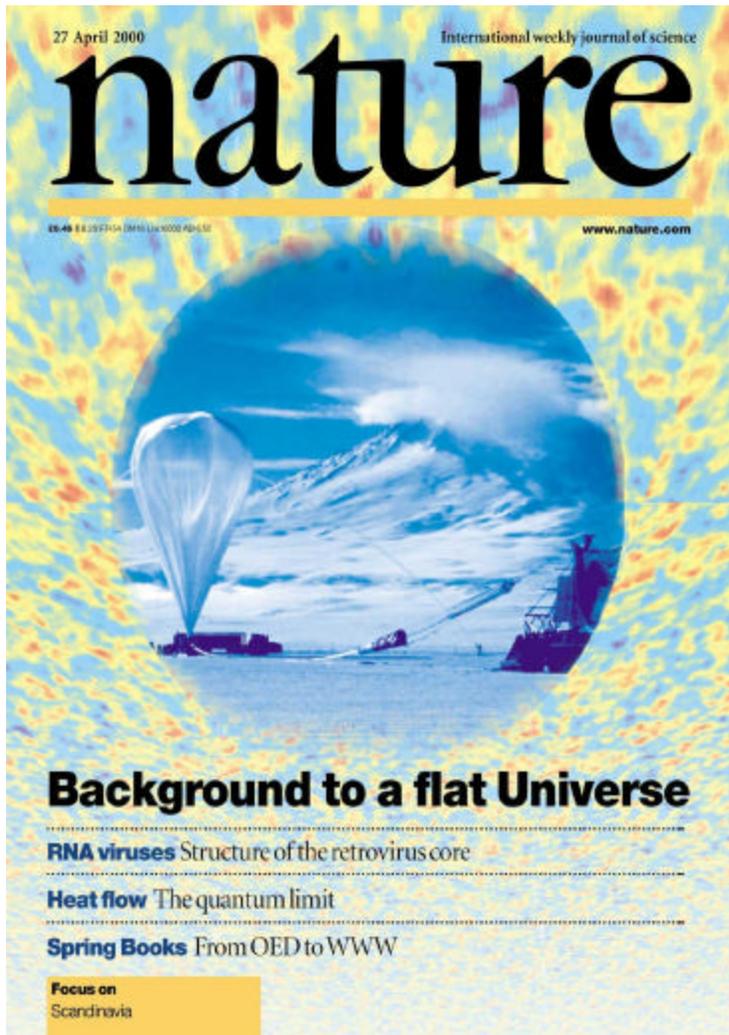
Science of Scale



<u>Project</u>	<u>Performance</u> (% of peak)
✍ Terascale Simulations of Supernovae	35%
✍ Accelerator Science and Simulation	25%
✍ Electromagnetic Wave-Plasma Interactions	68%
✍ Quantum Chromodynamics at High Temperatures	13%
✍ Cosmic Microwave Background Data Analysis	50%

Source: HDS' presentation to OMB examiner and to ASCAC advisory committee, 2003

Science of Scale: Cosmic Microwave Background Data Analysis



- ✍ **PI:** Julian Borrill, LBNL & UC Berkeley
- ✍ **Allocation Category:** Class B
- ✍ **Code:** Maximum likelihood angular power spectrum estimation (MADCAP)
- ✍ **Kernel:** ScaLAPACK
- ✍ **Performance:** 750 Mflop/s per processor (50% of peak)
- ✍ **Scalability:**
 - ✍ 0.78 Tflop/s on 1024 proc
 - ✍ 1.57 Tflop/s on 2048 proc
 - ✍ 3.02 Tflop/s on 4096 proc
- ✍ **Allocation:** 1.1 million MPP hours; requested and needs 2 million

Cosmic Microwave Background Data Analysis (cont.)



Recent accomplishments:

- ✍ MADCAP extended to enable simultaneous analysis of multiple datasets and CMB polarization – the new frontier.
- ✍ MADCAP was rewritten to exploit extremely large parallel systems, allowing near-perfect scaling from 256 to 4,096 processors.
- ✍ MADCAP++ is being developed using approximate methods to handle extremely large datasets for which matrix multiplications are impractical, such as will be generated by the PLANCK satellite.
- ✍ Recent results from NASA's WMAP satellite observations of the whole CMB sky confirm MADCAP analyses of previous partial-sky balloon datasets.

Status today



- ✍ We still get excellent science at NERSC -- but I have to work harder to explain it
- ✍ This challenge is everywhere
 - ✍ ASCI needs to relate platform performance to mission milestones
 - ✍ NAS study on the Future of Supercomputing
- ✍ As a community HPC needs to be more sophisticated to report about performance
- ✍ As a community HPC needs to show more unity in our communications

The Opportunity Today



National Coordination Office for
Information Technology Research and Development



The HECRTF seeks input about:

- (3) Federal Procurement of HEC Systems: This subtask will produce findings and recommendations that include:
- ✍ Identification of a strategy for developing practical performance measures for system procurement that correlate well with realized performance of actual applications
 - ✍ Recommended methods for deriving system performance targets from actual or projected application requirements or other user needs

see <http://www.hpcc.gov/hecrtf-outreach/>

The GC Question for the Future



What quantitative and objective measure of enabling scientific discovery can be used as the basis of evaluating computing facilities for basic science?