Workshop on High End Computing Revitalization

Dan Reed

reed@ncsa.uiuc.edu

Director, NCSA and the Alliance

Chief Architect, NSF ETF TeraGrid

William and Jane Marr Gutgsell Professor

University of Illinois

HEC Challenges

- **Today's HEC challenges …**
  - time to solution
    - too difficult to program and to optimize
    - better programming models/environments needed
  - often, efficiency declines with more processors
    - adversely affects time to solution and cost to solution
  - support overhead for system parallelism
    - management of large-scale concurrency
  - processor-memory latency and bandwidth
    - can be constraining for HEC applications
      - scatter-gather and global accesses
  - I/O and data management
    - volume and transfer rates
  - *power consumption, physical size and reliability*
- ***Have been around a long time …***

The most constant difficulty in contriving the engine has arisen from the desire to reduce the time in which the calculations were executed to the shortest which is possible.

Charles Babbage, 1791-1871

Many Workshops and Reports

- **Blueprint for Future Science Middleware and Grid Research and Infrastructure, August 2002**
  - http://www.nsf-middleware.org/MAGIC/default.htm
- **NSF Cyberinfrastructure Report, January 2003**
  - http://www.cise.nsf.gov/evnt/reports/toc.htm
- **DOE Science Network Meeting, June 2003**
  - http://gate.hep.anl.gov/may/ScienceNetworkingWorkshop/
- **DOE Science Computing Conference, June 2003**
  - http://www.doe-sci-comp.info
- **DOE Science Case for Large Scale Simulation, June 2003**
  - www.pnl.gov/scales/
- **DOE ASCR Strategic Planning Workshop, July 2003**
  - http://www.fp-mcs.anl.gov/ascr-july03spw
- ***Workshop on the Roadmap for the Revitalization of High End Computing, June 2003***
  - *www.cra.org/Activities/workshops/nitrd*
  - copies available at the back of the room

HECRTF Workshop Details

- ***Independent, community input* to HECRTF agencies**
  - HEC directions and needs
  - strategies and mechanisms
- **Strategic national needs/priorities**
  - discovery, competitiveness, defense and security
- **Community engagement**
  - collaborations, discussions and projects
- **Workshop participant charge**
  - same as that given to the government HECRTF group
- **Approach**
  - open call for white papers used to select participants
    - 84 white papers received
  - 220 workshop attendees on very little notice (weeks)

Working Groups, Chairs and Co-Chairs

1. **Enabling technologies**
   - Sheila Vaidya (LLNL) and Stu Feldman (IBM)
2. **HEC architecture – COTS-based**
   - Walt Brooks (NASA Ames) and Steve Reinhart (SGI)
3. **HEC architecture – Custom**
   - Peter Kogge (Notre Dame) and Thomas Sterling (Caltech/JPL)
4. **HEC runtime and operating system**
   - Rick Stevens (ANL) and Ron Brightwell (SNL)
5. **HEC programming environments and tools**
   - Dennis Gannon (Indiana) and Rich Hirsh (NSF)
6. **Performance modeling, metrics and specification**
   - David Bailey (LBL) and Allan Snavely (SDSC)
7. **Application-driven system requirements**
   - Mike Norman (UCSD) and John Van Rosendale (DOE)
8. **Procurement, accessibility and cost of ownership**
   - Frank Thames (NASA) and Jim Kasdorf (PSC)

Key Workshop Recommendations

- **Sustained investment**
  - research, development and system acquisition
    - key to long-term planning and strategic decisions
  - *see the virtuous cycle below*
- **Basic university research**
  - pipeline of ideas and people
    - attracting students and educating a new generation
  - research pipeline sustenance via stable funding
- **Deep collaboration**
  - academic researchers and government laboratories
  - industrial laboratories and computer vendors
  - *lower the barriers for collaboration/technology transfer*
- **Multiple iterations of the virtuous cycle**
  - advanced research and development
  - large-scale system prototyping
  - product development and assessment
  - *deploy, learn, deploy, learn, deploy …*

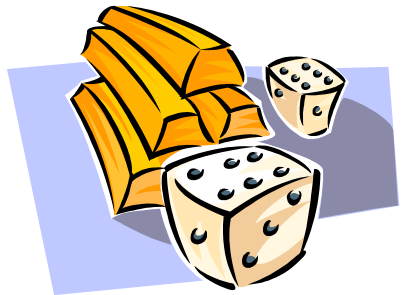

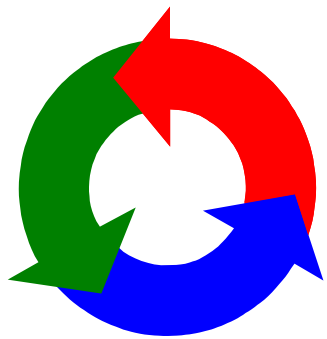

Workshop Recommendations

- **Enabling technologies**
  - power management and interconnection performance
    - new device technologies, three-dimensional integration and packaging
  - long-term research in novel devices
    - superconducting technologies and spintronics
    - photonic switching and molecular electronics
  - software for large scale systems
    - scaling demonstrations and reduced time to solution
    - real-time performance monitoring and feedback
- **COTS technology trends**
  - memory-class ports for high bandwidth, low latency interconnects
  - higher-speed signaling and higher radix routers
  - field programmable gate arrays (FPGAs)
- ***Our force multiplier is early research and development***
  - *product development is very late*

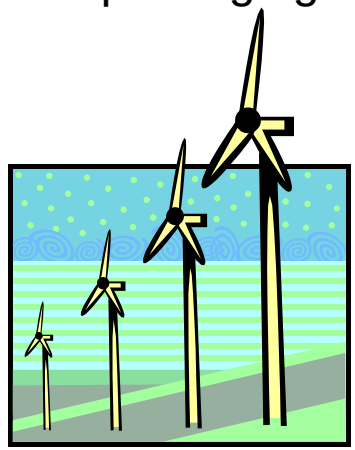

Workshop Recommendations

- **Custom architectures**
  - architectural approaches
    - spatially direct mapped, vectors and streaming architecture
    - processor in memory architecture and special purpose devices
  - dynamic resource management software
  - programming models that explore algorithmic parallelism
  - proof of concept assessment

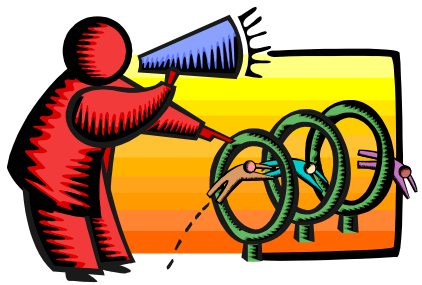

- **Runtime and operating systems**
  - alternate resource models
    - performance feedback for dynamic adaptation
    - increased coupling among operating system, runtime and applications
    - new models for I/O coordination and security
  - revolutionary, rather than evolutionary system software research
  - investment in large-scale testbeds

Workshop Recommendations

- **Performance analysis**
  - reduced time to solution for applications
  - coordinated benchmarking for cross-vendor use
  - enhanced performance modeling, monitoring and analysis

- **Programming environments and tools**
  - support for multidisciplinary, multiscale applications
  - higher investment in the quality, availability and usability of tools
  - interoperable libraries and software
  - structural changes in funding approaches
    - software capitalization program
    - institute for software development

Software

- **Hardware and software both matter**
  - hardware is necessary but not sufficient
    - like a car without a road
  - look at our experiences
    - CDC 6600, Cray 1, CM-5, KSR, …

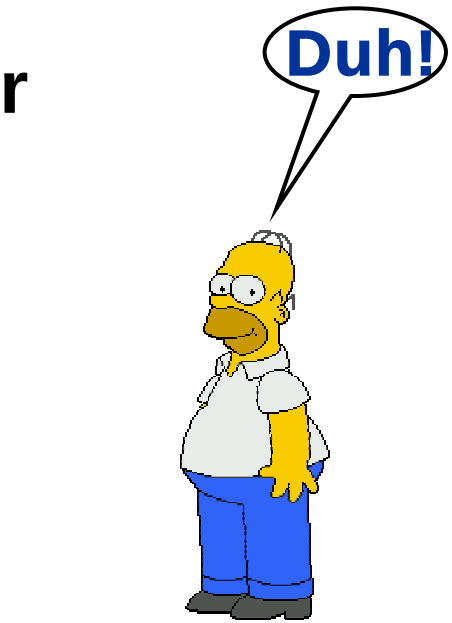

- **Software is often overlooked**
  - both application and infrastructure
  - application embodies the science
  - tools enable/hinder productivity
  - it is not cheap!

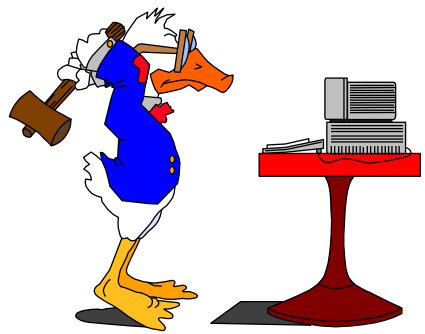

Workshop Recommendations

- **System procurement**
  - functional specifications to define science requirements
  - total cost of ownership as primary evaluation criteria
  - collaborative procurement strategies

- **Applications: *the reason for this!***
  - dramatically enhanced system capabilities
    - sustained performance of 20-100 TF
  - multidisciplinary teams
    - application and computer scientists

Defense and Security Challenges

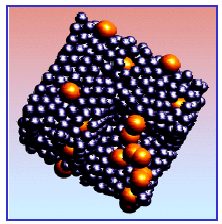

Basic Research

Simulating High-Energy Density Rocket Fuels

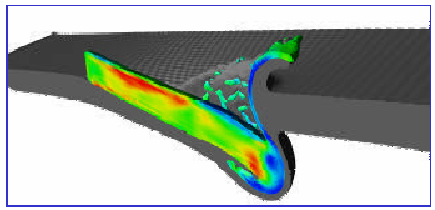

Advanced Technology

Armor and Projectile Design

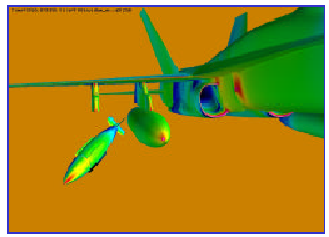

Developmental T&E

Support of Aircraft-Store Compatibility and Weapons Integration

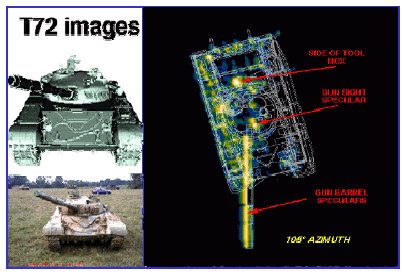

Intelligence

Radar Cross-Sections Predictions

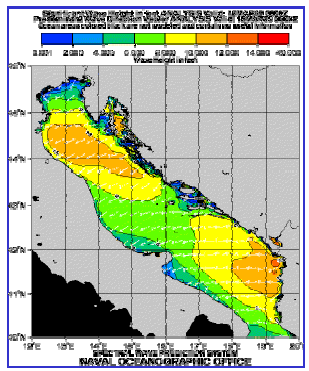

Operations

Ocean/wave forecasting

- Nuclear stockpile stewardship
- Cryptography and climate data assimilation
- Reusable launch vehicle design
- Survivability/stealth, engineering design

Source: John Grosh/David Nelson

Physics/Cosmology Challenges

- **The Standard Model**
  - quantum theory
    - electroweak and strong forces
- **Gravity and relativity**
  - no “Grand Unified Theory”
    - gravity integration and rationale for mass
- **Dark matter and dark energy**
  - most of the universe’s mass is “invisible”
    - mass candidates: MACHOS (baryonic), WIMPS, …
  - expansion seems to be accelerating
    - vacuum energy (cosmological constant)
    - dynamic field (quintessence)
- **Recent experimental data**
  - Wilkinson Microwave Anisotropy Probe
    - universe is 13,400 ± 300 million years old and flat

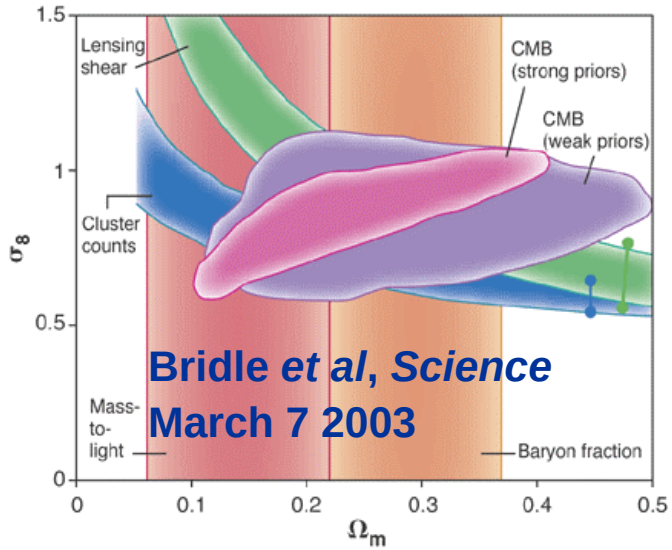

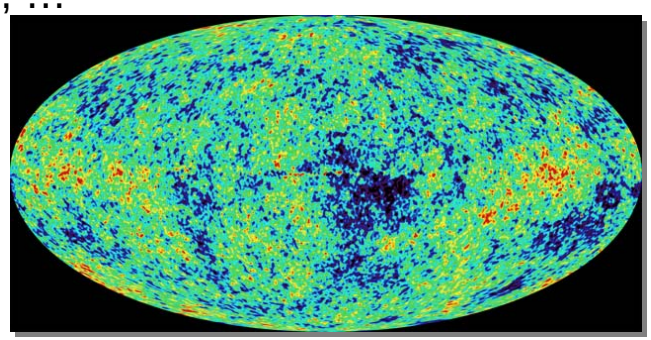

Biology Challenges

DNA sequence

Sequence Annotation

Protein sequence and regulation

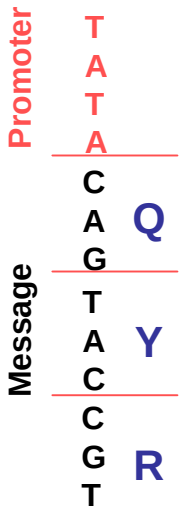

Homology based protein structure prediction

Protein structure

Molecular simulations

Protein/enzyme function

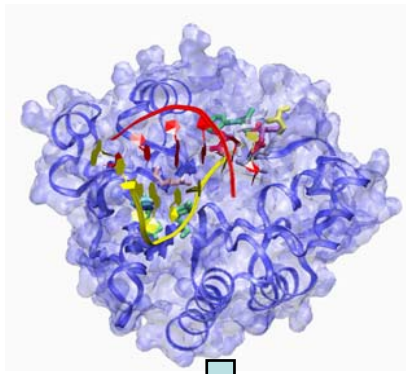

Data integration

Multi-protein machines

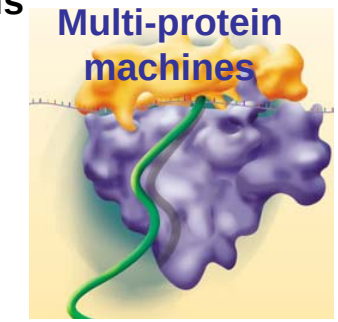

Network analysis

Pathway simulations

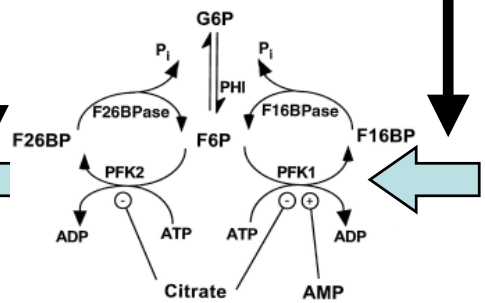

Metabolic pathways and regulatory networks

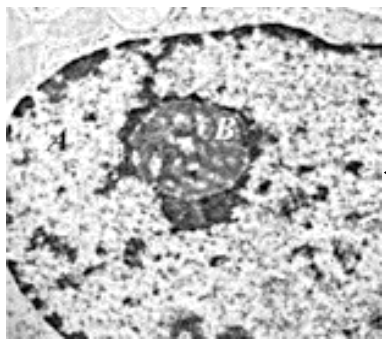

Bacteria and cells

Organs, Organisms and Ecologies

NCSA

Strategic Planning *Really* Matters

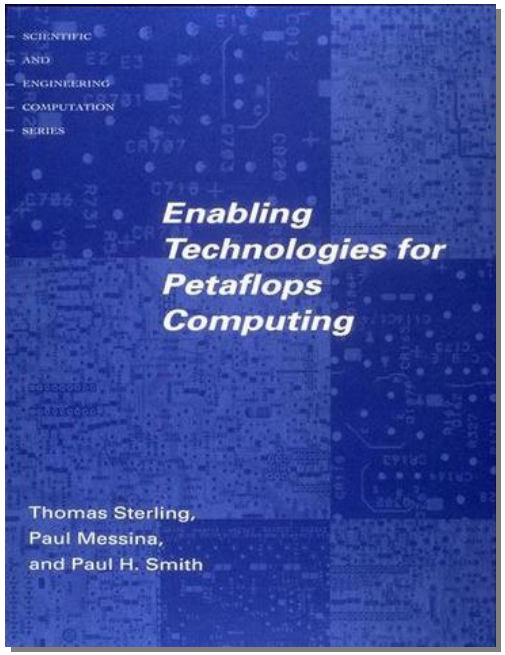

February 1994

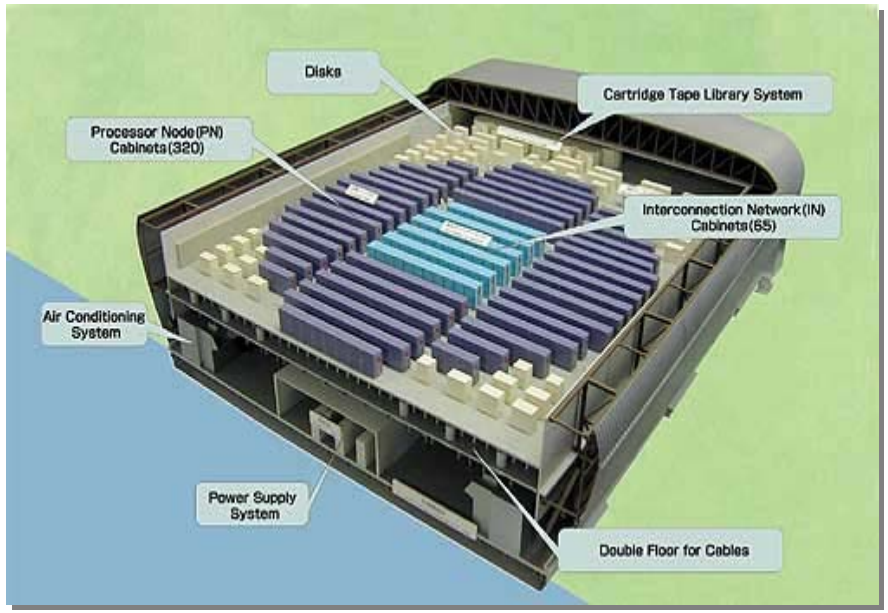

- **Earth System Simulator**
  - long-term decadal plan
  - multiple hardware generations and resource commitments

HEC: The Integrated View

- **Go long**
  - sustained, long-term exploration
    - long-term, strategic plans and support
    - coupled with tactical evolution and exploitation
  - 10-15 years is about the right time scale
    - 2-3 iterations to get it right

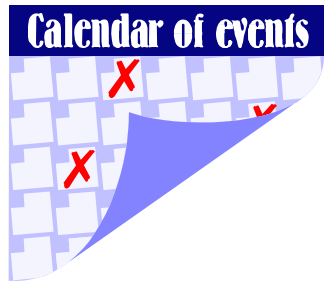

- **Go deep**
  - multidisciplinary examination of problems
    - understand applications and needs
    - architecture, systems, software and algorithms
  - resource scale matched to problem scale

- ***Remember and act***
  - *1999 PITAC report recommendations*
    - *not yet fully implemented*

The Cambrian Explosion

- **Most phyla appear**
  - sponges, archaeocyathids, brachiopods
  - trilobites, primitive mollusks, echinoderms
- **Indeed, most appeared quickly!**
  - Tommotian and Atdbanian
  - as little as five million years
- **Lessons for computing**
  - it doesn't take long when conditions are right
    - raw materials and environment
  - leave fossil records if you want to be remembered!