HECRTF, PITAC and the Future

Dan Reed

Director, RENCI

William R. Kenan, Jr. Eminent Professor

Dan_Reed@unc.edu

University of North Carolina at Chapel Hill

Duke University

North Carolina State University

How Big Is Big?

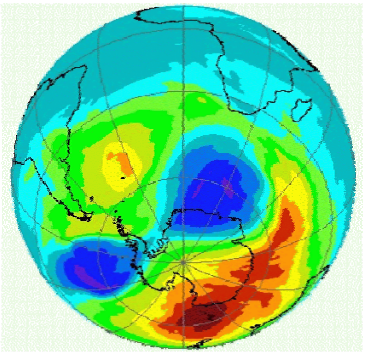

- **Every 10X brings new challenges**
  - 64 processors was once considered large
    - it's now a research cluster in a closet
  - 1024 processors is today's "medium" size
  - 2048-8096 processors is today's "large"
    - we're struggling even here
- **10K-100K processors is in sight**
  - we have fundamental challenges …
  - … and no integrated research program

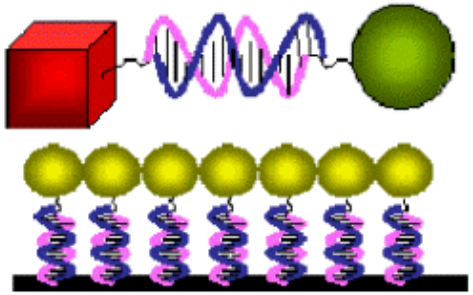

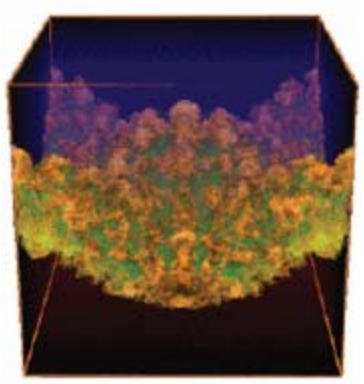

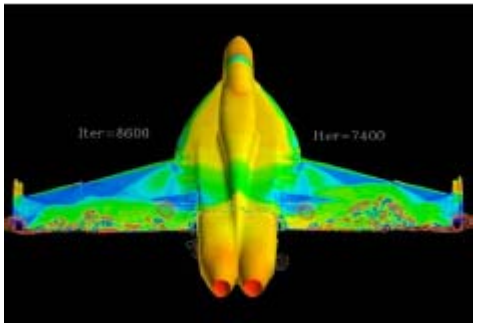

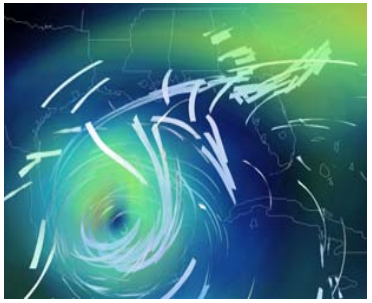

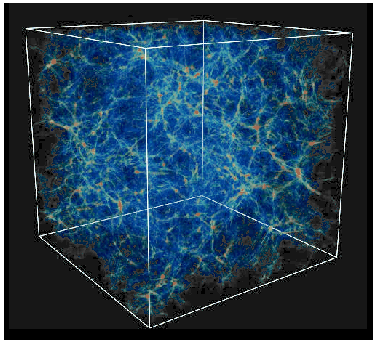

The Siren Call: Peak Performance

The Sirens inhabited an island surrounded by dangerous rocks. They sang so enchantingly that all who heard were drawn near and shipwrecked. Jason and the Argonauts were saved from them by the music of Orpheus, whose songs were lovelier. Odysseus escaped them by having himself tied securely to a mast and by stopping the ears of his men.

The Siren Call: Peak Performance

- **1890-1945**
  - mechanical, relay
  - 7 year doubling
- **1945-1985**
  - tube, transistor,..
  - 2.3 year doubling
- **1985-2003**
  - microprocessor
  - 1 year doubling
- **Every year**
  - equal to all previous history
- **Storage, networks and graphics**
  - even faster!
- ***Delivered performance and software development***
  - *dependent on algorithms and architecture match*
  - *a much more nuanced (and often ugly) story …*

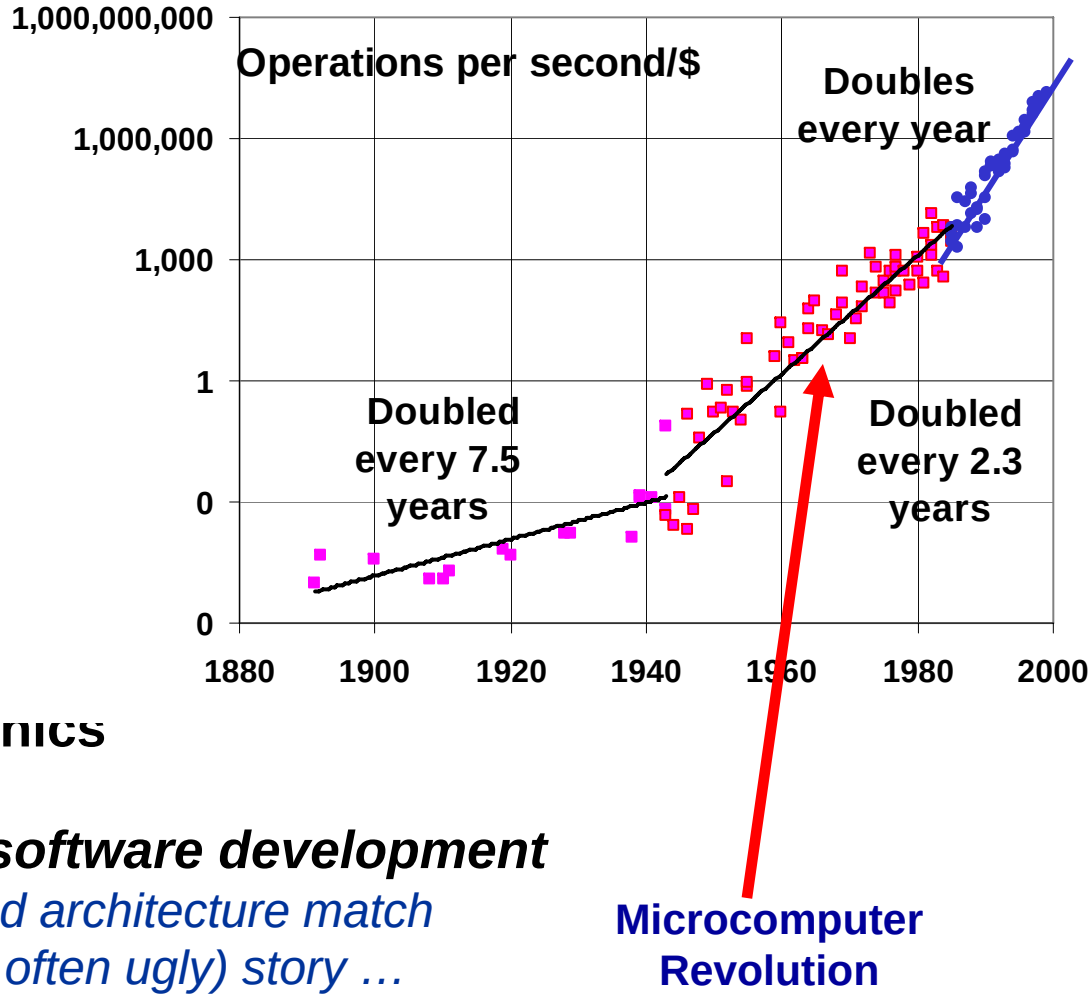

Data source: Jim Gray

High-End Computing 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
  - fraction of single processor peak is very low (5-10%)
- **Support overhead for system parallelism**
  - management of large-scale concurrency
- **Processor-memory latency and bandwidth**
  - constraining for HEC applications
    - scatter-gather and global accesses
- **I/O and data management**
  - volume and transfer rates
- **Power consumption, physical size and reliability**

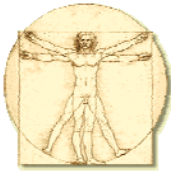

Many Workshops and Reports

- **Computation as a Tool for Discovery in Physics, S**
  - www.nsf.gov/pubs/2002/nsf02176/start.htm
- **Blueprint for Future Science Middleware and Grid August 2002**
  - www.nsf-middleware.org/MAGIC/default.htm
- **NSF Cyberinfrastructure Report, January 2003**
  - www.cise.nsf.gov/evnt/reports/toc.htm
- **DOE Science Network Meeting, June 2003**
  - gate.hep.anl.gov/may/ScienceNetworkingWorkshop/
- **DOE Science Computing Conference, June 2003**
  - www.doe-sci-comp.info
- **DOE Science Case for Large Scale Simulation, Ju**
  - www.pnl.gov/scales/
- **DOE ASCR Strategic Planning Workshop, July 20**
  - www.fp-mcs.anl.gov/ascr-july03spw
- **Roadmap for the Revitalization of High End Comp**
  - www.hpcc.gov/hecrtf-outreach
- **House Science Committee Hearing, "Supercompu**
  - **www.house.gov/science/hearings/full03/index.htm**
- ***PITAC Computational Science, 2004-2005***
  - *stay tuned*

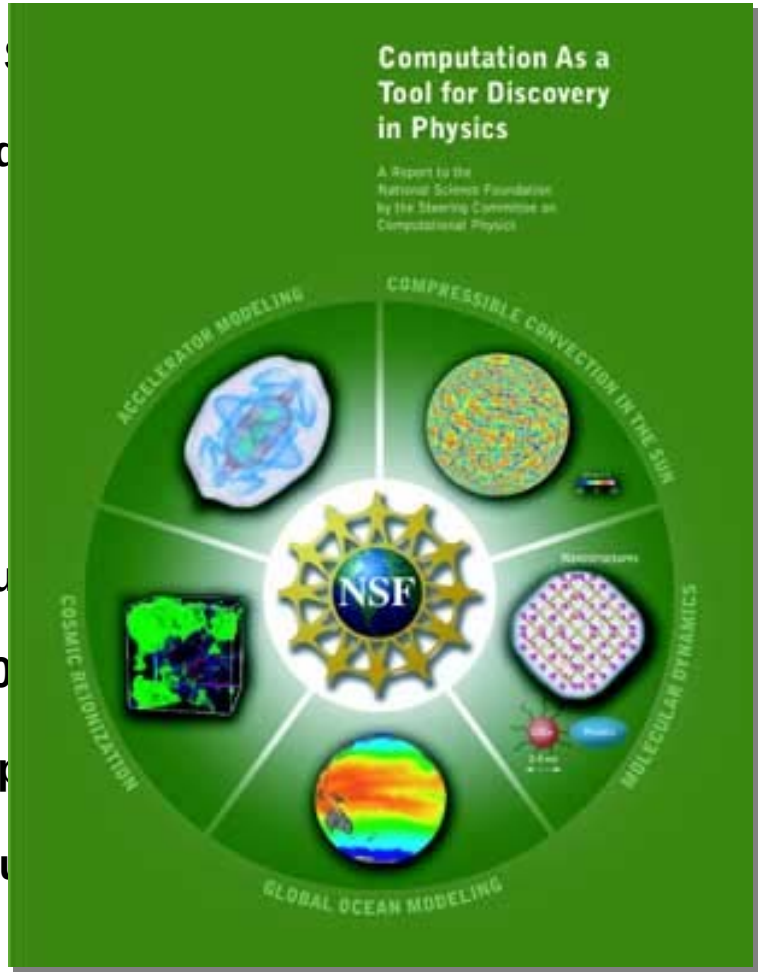

FY 2003 Federal Budget

"Due to its impact on a wide range of federal agency missions ranging from national security and defense to basic science, high end computing—or supercomputing —capability is becoming increasingly critical. Through the course of 2003, agencies involved in developing or using high end computing will be engaged in planning activities to guide future investments in this area, coordinated through the NSTC. The activities will include the development of interagency R&D roadmap for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues (along with recommendations where applicable) relating to federal procurement of high-end computing systems. The knowledge gained for this process will be used to guide future investments in this area. Research and software to support high end computing will provide a foundation for future federal R&D by improving the effectiveness of core technologies on which next-generation high-end computing systems will rely."

HECRTF

- **High End Computing Revitalization Task Force (HECRTF)**
  - requested by Congress in FY03 budget language
  - commissioned by Office of Science and Technology Policy (OSTP)
  - coordinated through National Science and Technology Council (NSTC)
- **Charge**
  - develop *a five year plan* to guide future Federal HEC investments
    - overall strategy for HEC investments for FY05-FY09
- **Mechanisms**
  - participation by agencies using/developing HEC systems
  - multiple interagency working groups
    - task integration
    - core technologies research and development
    - capability, capacity and accessibility
    - procurement of federal HEC systems
  - input to FY2005 federal budget
  - coordination by NITRD office
- ***Roadmap for the Revitalization of High End Computing, June 2003***
  - *www.hpcc.gov/hecrtf-outreach*

HECRTF Chairs

- **Co-chairs**
  - John Grosh (DoD)
  - Alan Laub (DOE/SC)
- **Task chairs**
  - HEC core technologies R&D
    - Kamal Abdali (NSF)
    - José Muñoz (DOE/NSSA)
  - capability, capacity and accessibility
    - Judy Devaney (NIST)
    - Tom Page (NSA)
  - procurement
    - Bill Turnbull (NOAA)
    - Phil Webster (NASA)
- **Community workshop chair**
  - Dan Reed

Interagency Perspectives*

- **HEC is a declining fraction of the overall market**
  - future systems may be less suitable to HEC needs
- **Future success will require coordinated effort**
  - R&D and engineering of new architectures and systems
  - software research and development
    - systems and middleware
    - programming environments and applications
  - new domain science and algorithms
  - procurement of new COTS and custom systems
    - sustainable strategies
- **Targeted funding of HEC systems may be required**
    - including development of new systems

- ***My assessment; my apologies for any misrepresentations***

But It's Not A New Problem …

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

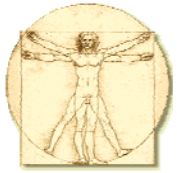

HECRTF Schedule (2003)

- **February 28**
  - memorandum announcing HECRTF
- **March 10**
  - kick-off meeting
- **April 18**
  - call for white papers on HECRTF charge
- ***June 16-18***
  - *CRA workshop with academia, industry and government*
- **July 14, 21**
  - meetings with industry under non-disclosure
- **August 29**
  - plan submitted to OSTP (Office of Science and Technology Policy)

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)

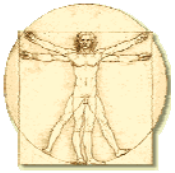

Working Groups, Chairs and Co-Chairs

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)

HECRTF Recommendations

- **Sustained investment**
  - research, development and system acquisition
    - key to long-term planning and strategic decisions
  - *see the virtuous cycle (stay tuned)*
- **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 …*

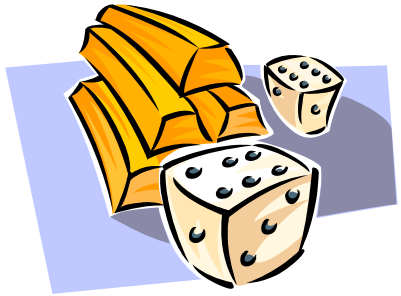

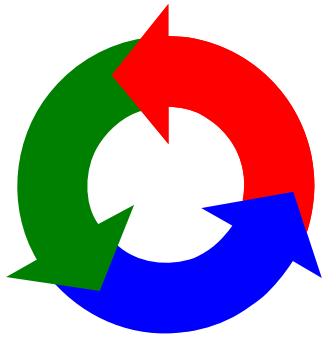

Workshop Recommendations

- **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

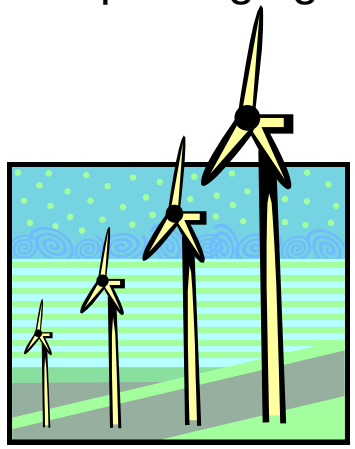

- **COTS technology trends**
  - memory-class ports for high bandwidth, lower latency interconnects
  - higher-speed signaling and higher radix routers
  - field programmable gate arrays (FPGAs)
- ***Our force multiplier is early research and development***
  - *product development is very late*
  - *don't confuse momentum with progress!*

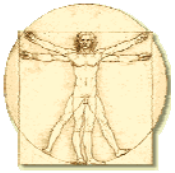

Workshop Recommendations

- **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

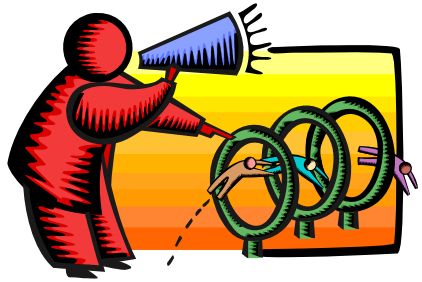

- **Runtime and operating systems**
  - alternate resource models
    - performance feedback for dynamic adaptation
    - increased coupling among operating system, runtime and applications
    - new models for I/O coordination and security
  - revolutionary, rather than evolutionary system software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

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!

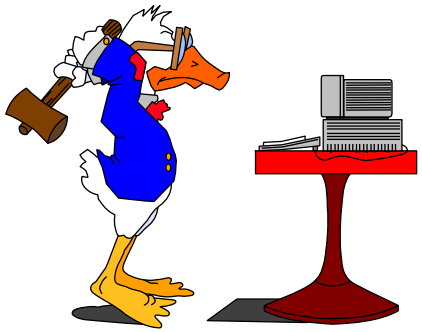

Observations on Software

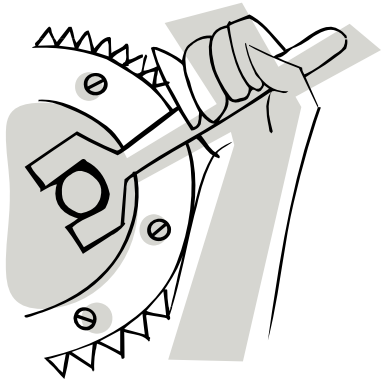

- **Business**
  - capital is cheap
  - labor is expensive
  - costs are usually explicit
    - and had better be lower than revenues!
- **Academia and government**
  - capital is (seemingly) expensive
  - labor is (seemingly) cheap
    - student and faculty time
  - costs are usually implicit
    - and often skew realistic assessment

- **This is a critical issue with respect to software**
  - 10 one-year programs are not equivalent to one ten-year program

Software Complexity and Growth

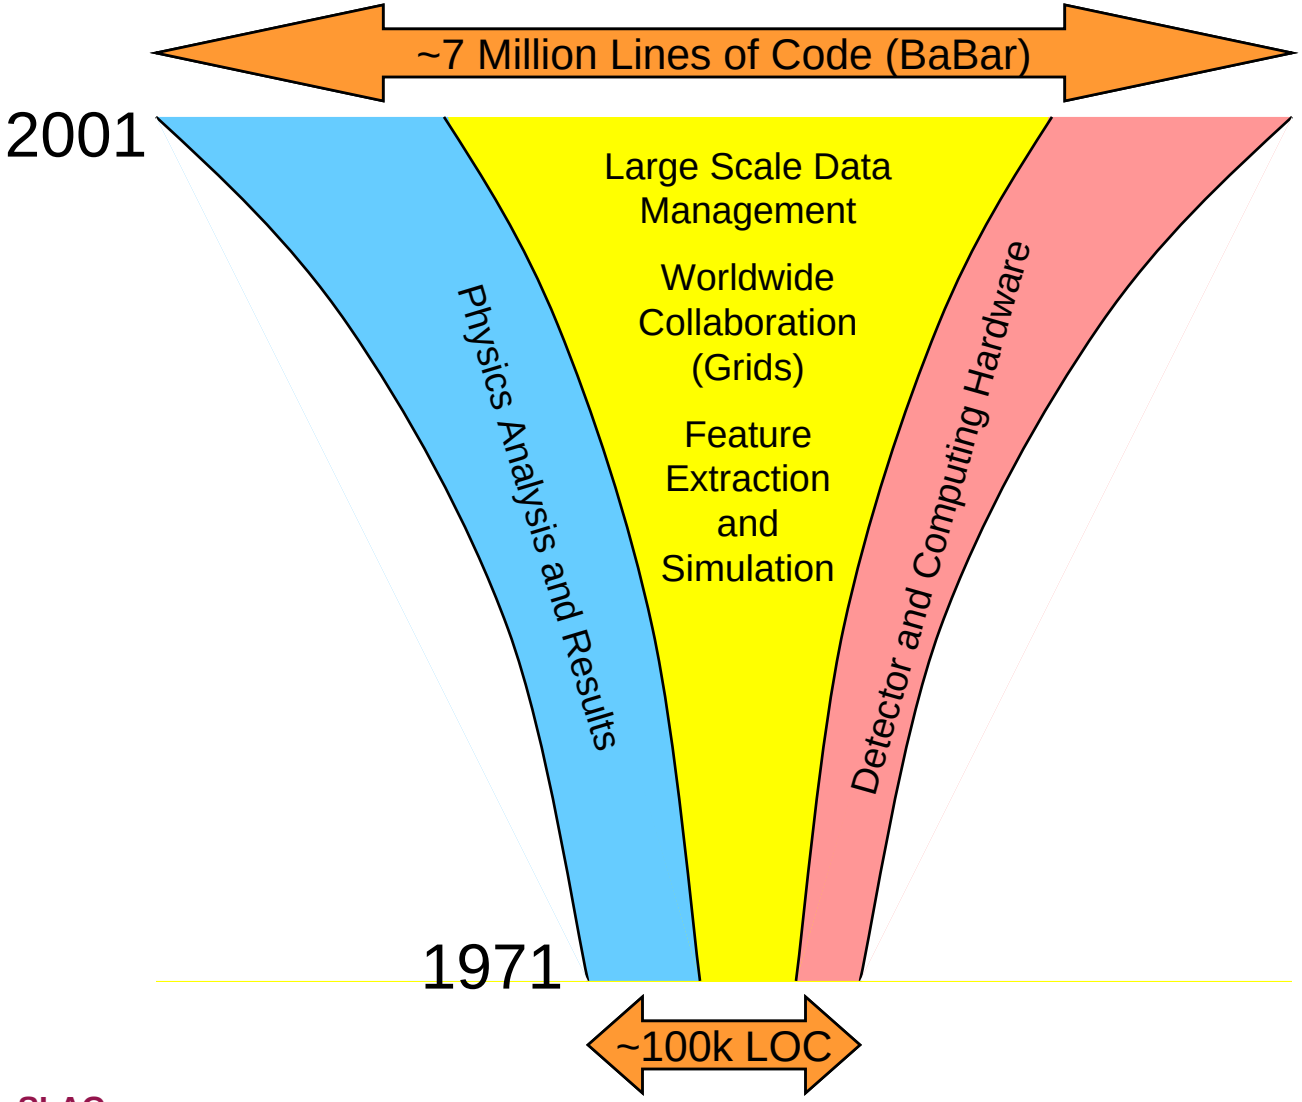

Source: Mount, SLAC

Workshop Recommendations

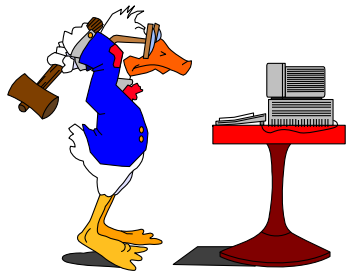

- **Performance analysis**
  - reduced time to solution for applications
  - coordinated benchmarking for cross-vendor use
  - enhanced performance modeling, monitoring and analysis

- **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

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

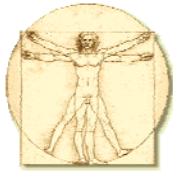

PITAC and Computational Science

- **Three PITAC sub-committees**
  - health (well underway)
  - security (launching)
  - science and engineering (just starting)
    - Dan Reed (chair)
- **We will look broadly at IT and science**
  - research, infrastructure and community
- **House Science Committee hearing**
  - April 29 on NITRD reauthorization and HPC
- ***I need your input***
  - *challenges and opportunities*

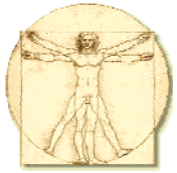

The Cambrian Explosion

- **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!