Disclaimer

These slides were presented at the President's Information Technology Advisory Committee's (PITAC) November 4, 1998 meeting by the chairs of its six panels. The panels were asked to suggest revisions to the PITAC's Interim Report. The information in these slides will be taken into consideration as the PITAC drafts its final report.

High-End Computing

PITAC SUB-COMMITTEE

November 4, 1998

High-End Committee

- ❶ Innovative Computing Technologies and Architecture
- ❷ Software for improving the performance of high-end computing
- ❸ Sustain petaflop/petaop on real applications
- ❹ Acquire the most powerful computing systems to support science and engineering

Format of Presentation

- Analysis used for recommendation
- Low, Medium, and High Level of incremental funding. (relative to current level in absolute dollars)
- Begins with FY2000 funding, for 5 years

PITAC High-End Sub-panel

Innovative Technologies & Architectures

George Cotter   David Patterson

Thomas Sterling   Mary Vernon

Findings

- Breakthroughs in sustained performance and ease of use will require innovations in architecture and component technologies
- architecture and technology research are expensive and time consuming
- capability exists to carry out research of a broader scope and at an increased pace

Refinements to the Recommendation

- Architectures for high-end computing must be derived in concert and iteratively with hardware and software technology research
- New initiatives should include investigation of alternative architectures and technologies that are of high risk & possibly high payoff
- New investment is needed in
  - improved simulation & analysis tools
  - full system architectural design experiments
  - device technology prototype fabrication facilities

New Initiatives: Low Threshold

- Program Scope:
  - detailed design studies/simulation of alternative high-end architectures
  - detailed simulation of multiple possible device types, and carry promising device to required laboratory quality implementation

New Initiatives: Medium Threshold

- Program Scope:
  - prototype implementation of critical elements of future generation high-end computer architectures
  - expand the set of advanced technology types for experimental fabrication; develop an advanced technology for full-scale implementation

New Initiatives: High Threshold

- Program Scope:
  - implement a full system with advanced architecture, component technology & software; full-scale system simulation for software development
  - implement a new fabrication line and process for a particularly promising new device/technology

PITAC High-End Sub-panel Software

Jack Dongarra
Dan Reed, George Spix
Irwing Wladawsky-Berger

“Software” Broadly Means ...

- Algorithms, methods, and libraries
- Compilers, tools, and runtime systems
- Operating systems and resource management
- Problem solving environments
- Distributed computing and collaboration
- Visualization and data management

Background and Observations

- High-end software
  - an old (and worsening problem)
  - fragile and ill-matched to emerging hardware
  - effectively supports only the “hardy pioneers”
  - quality/quantity limited by funding
- Good software takes time and money
  - support for common idioms
  - “get it wrong” before you “get it right”

High-end Software Perspectives

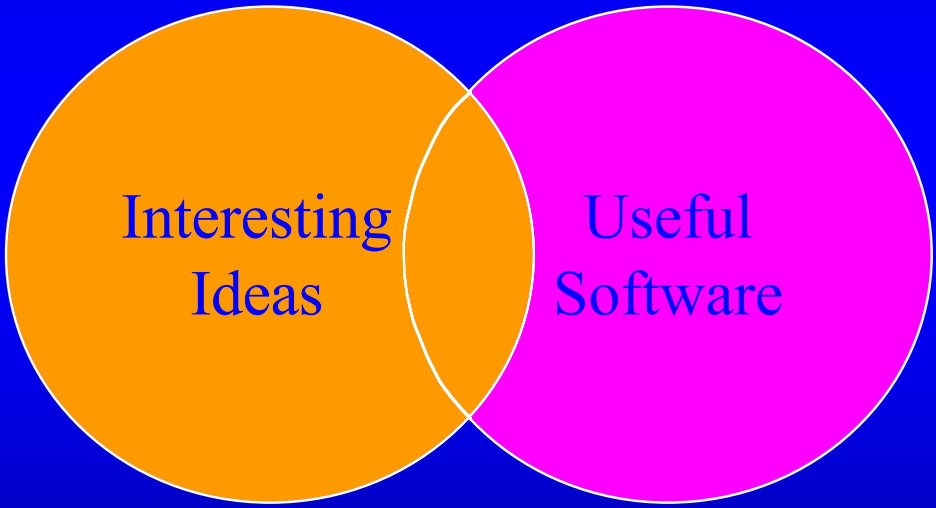

Emerging Challenges

- System complexity
  - thousands of interconnected microprocessors
  - deep, multilevel memory hierarchies
  - hundreds of secondary/tertiary storage devices
  - high-resolution visualization hardware
  - high-speed, wide-area networks
- Application complexity
  - distributed computational grids
  - multi-lingual, multidisciplinary applications

Observations

- Transition from research to practice takes time
  - must look to the long term
- Software has been *dramatically* underfunded
- This has exacerbated our current situation
  - experimental systems with unstable software
  - highly variable achieved performance
- We need to act now to prevent a recurrence

Recommendations

- At least five targeted areas
  - scalable algorithms and libraries
  - integrated programming models/tools
  - intelligent data management
  - scalable visualization and steering (HCI)
  - end-to-end assessment
- Funding models
  - single investigator and modest groups
  - software centers and expedition centers

Software Funding Paucity

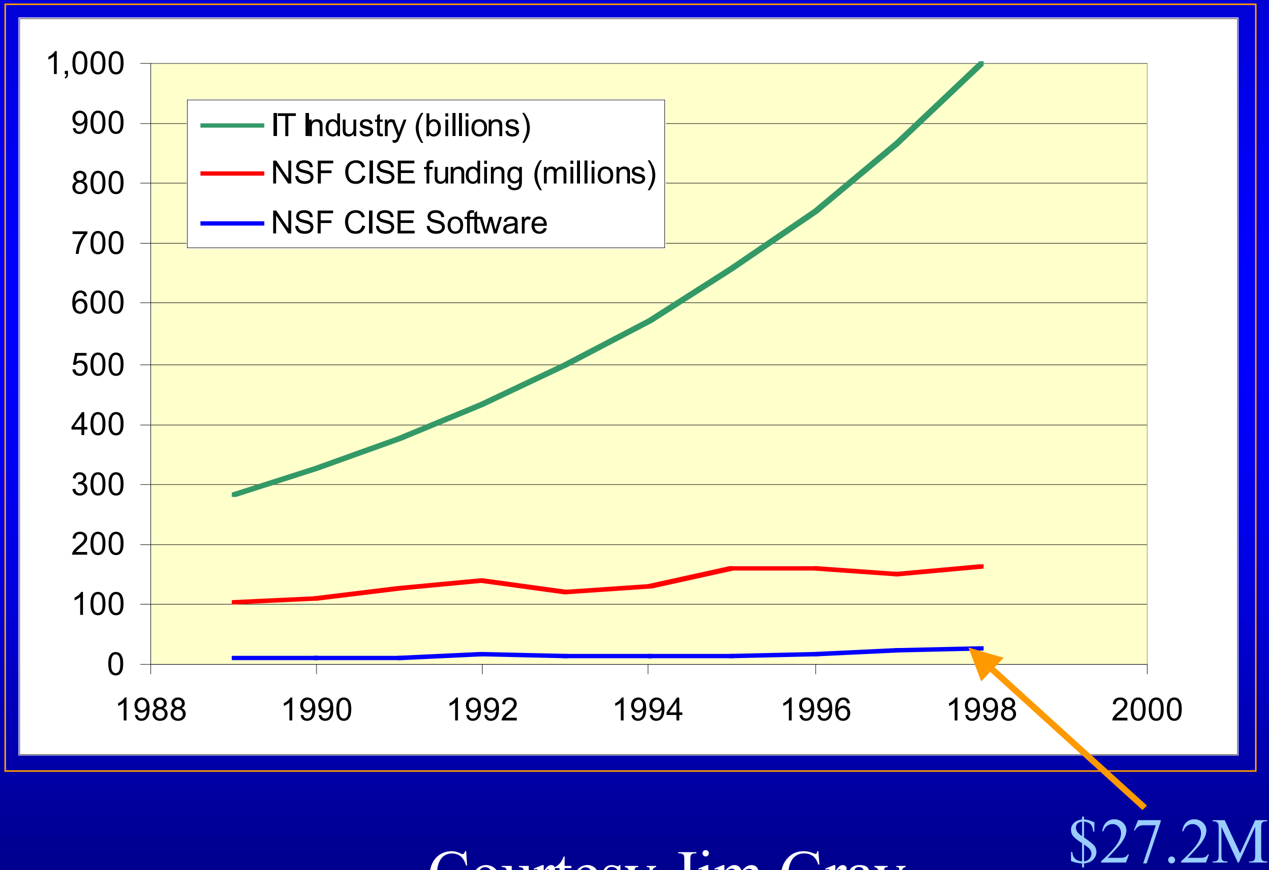

Courtesy Jim Gray

PITAC High-end Sub-panel
Sustain Petaops/Petaflops

George Cotter

Thomas Sterling

Steve Wallach

Topics

- Summary of Interim Recommendation
- Technical Issues
- Advanced Peta(fl)ops Program Objectives
- Proposed Peta(fl)ops Program Tasks
  - 5 year plan, FY00 start
  - low, medium, & high funding levels
- Budget Breakdown

Summary of Interim Findings and Recommendations

- Major technological advances required for peta(fl)ops by 2010.
- Technology research driver in devices, architecture, software, and algorithms.
- Increase funding of HECC Peta(fl)ops Program.
- Requires a well balanced effort on software and hardware.
- Technical comparison of current architectures with promising new ones.

Technical Issues

- Peta(fl)ops architecture archetypes
  - clustered systems
  - tightly coupled multiprocessors
  - hybrid technology, multithreaded
  - processor in memory
  - special purpose devices
  - metacomputing
- Distinguishing Factors
  - multi million-way parallelism
  - very long latencies of 10,000 cycles or more
  - bi-section bandwidths of 10^{15} bytes per second
  - 100X power efficiency improvement
  - 50 to 1000 Tbyte main memory

Technical Issues (continued)

- Very high level of system software complexity to manage clusters of > 100K processors
- SPDs of limited utility for general computing although a technical driver.
- IPG uncertain for wide-usage petaflops computation.
- Today's hardware technology, system software, and algorithms inadequate
- Possible that conventional MPP paradigm is inappropriate for 1000X performance gain

Strategic Objectives

- Establish requirements, operational boundaries, basic research needs, and alternative structures for peta(fl)ops scale computing systems and applications.
- Provide driver goals for recommended programs in architecture/technologies and system software.
- Determine feasibility of rapid path to peta(fl)ops performance in 2004 to 2007 timeframe.
- Enable new choices to future US directions in high end computing through investment in critical device fabrication facilities and pathfinding prototypes.
- Establish a single-site of coordination.

Detailed Objectives

- Characterize critical applications and assess requirements.
- Develop important architectural approaches with emphasis on latency/resource management, efficiency, and feasibility of implementation.
- Develop innovative software methodologies capable of managing $>10^5$ processors, $>10^4$ latency, $>10^{15}$ bandwidth, and $>10^{14}$ memory.
- Advance inchoate device technologies exhibiting high potential for accelerating achievement of peta(fl)ops.
- Derive algorithms exhibiting latency tolerance and yielding very high parallelism.

Proposed Program Tasks Low Funding Level

Establish base foundation of quantitative knowledge determining petaflops opportunity space plus focus activity.

- Scaling studies of algorithms to determine requirements.
- Identify petaflops scale applications critical to national needs and objectives.
- Roadmaps for candidate technologies.
- Implications of petaflops to system software functionality.
- Experimental testbed as focus and risk reduction for high potential alternative approach.
- Annual report on state of understanding on Pflops.

Proposed Program Tasks Medium Funding Level

Set vector to petaflops with focus project.

- Detailed design studies of petaflops architectures.
- System software approach for each architecture.
- Device technology design and analysis coupled with specific architecture designs.
- New scalable algorithms for important kernels.
- Augmentation of pathfinding testbed.

Proposed Program Tasks
High Funding Level

Go for it.

- Detailed simulation studies of candidate architectures.
- Prototype fabrication facilities for critical technologies.
- Performance driven O/S, compilers, and tools for target architecture.
- Implementation of scalable mission-driven applications.
- Petaflops computer implementation engineering studies by collaborative teams.
- 2nd prototype testbed.
- Initiate design of planned petaflops computer.

A Hypercomputer for the 21st Century

1 Petaflops

1 Petabyte

1 Pbyte/second

2.6 Mwatts

FY2005

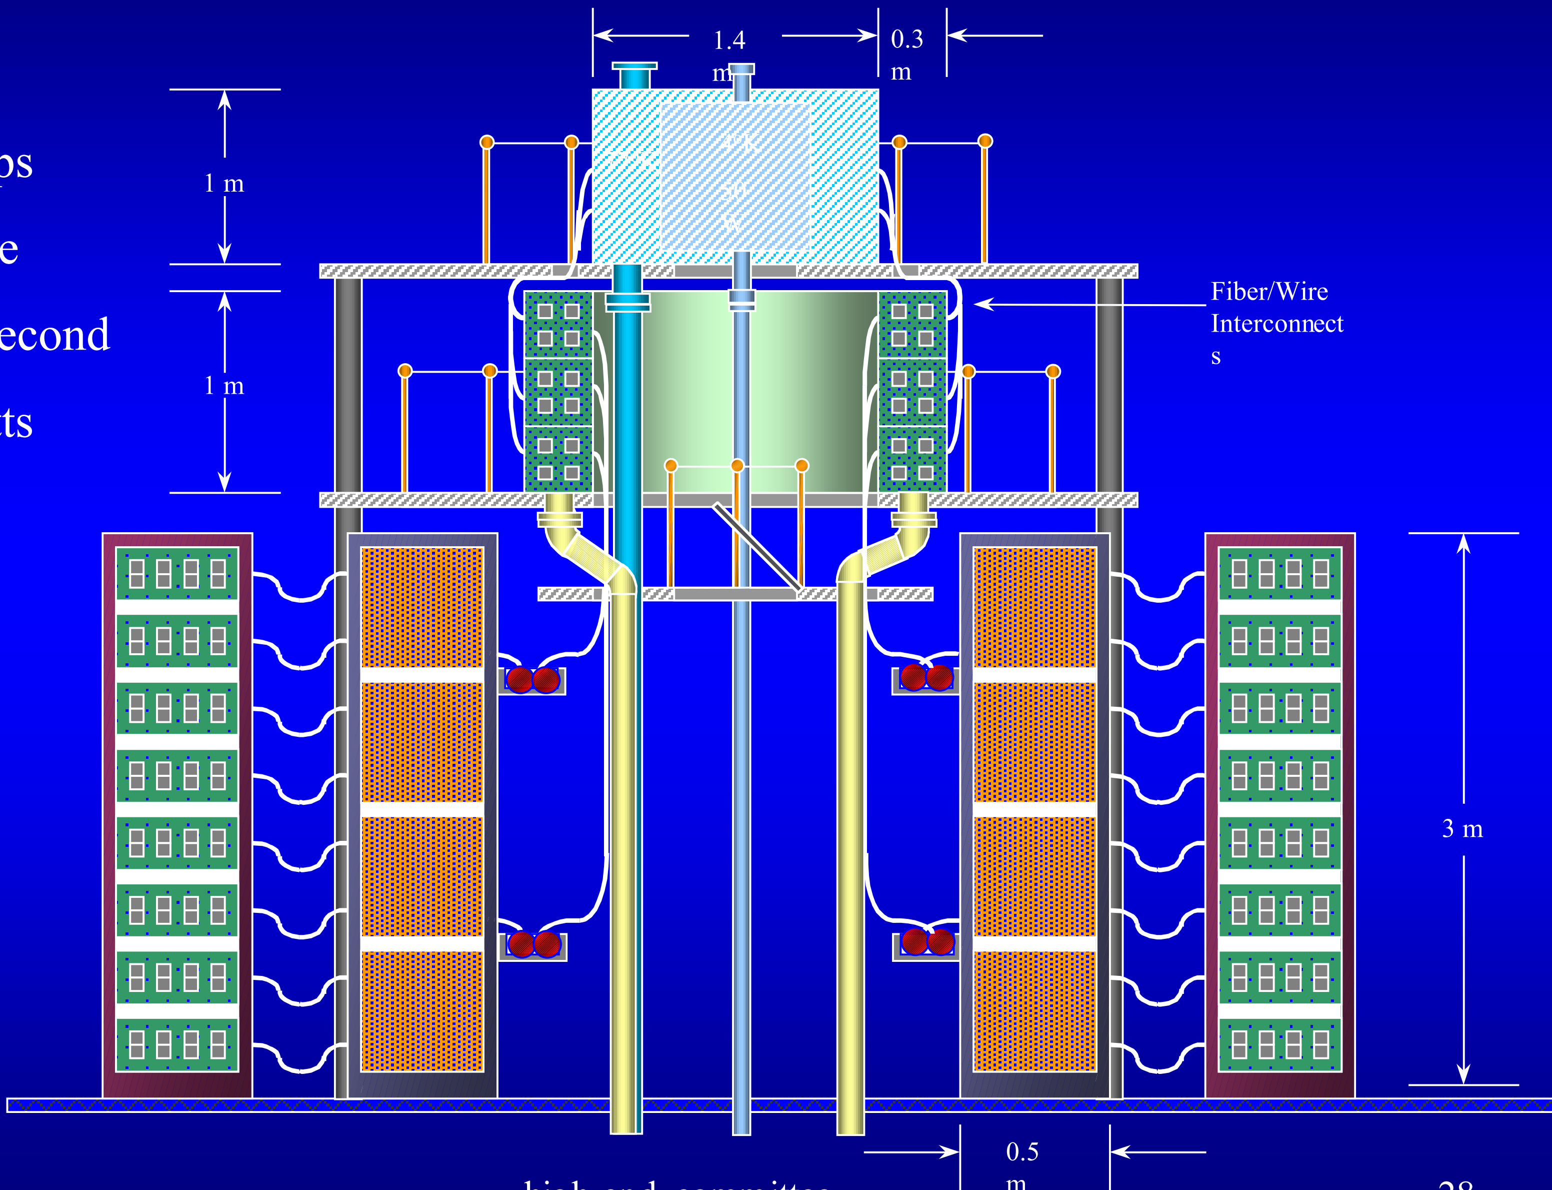

PITAC High-End Sub-panel Acquisition of High-End Systems

Steve Wallach Larry Smarr
Dave Cooper Bo Ewald

Strategic Objective

- Fund the acquisition of the most powerful high-end computing systems to support science and engineering research

Assumptions and Observations

- IT R&D is divided into 2 areas
  - Development of the IT technology (e.g., architectures, hardware, software, etc.)
  - Researchers that use high-end systems to support their research (e.g., computational physics|chemistry|biology|mechanics)

Assumptions and Observations

- Several government agencies and their users are falling behind the power curve.
  - NSF
  - DOE/ER
  - NIH
  - NOAA
  - NASA
  - EPA
- DOD and DOE (weapons) have their own high-end acquisition budgets

Grouping of Users

❶ Requires immediate action
  – NSF (PACI Centers)

❷ Requires a planning document within 6 months
  – DOE/ER, NIH, NOAA, NASA, EPA

High End Capacity Available to National Academic Researchers

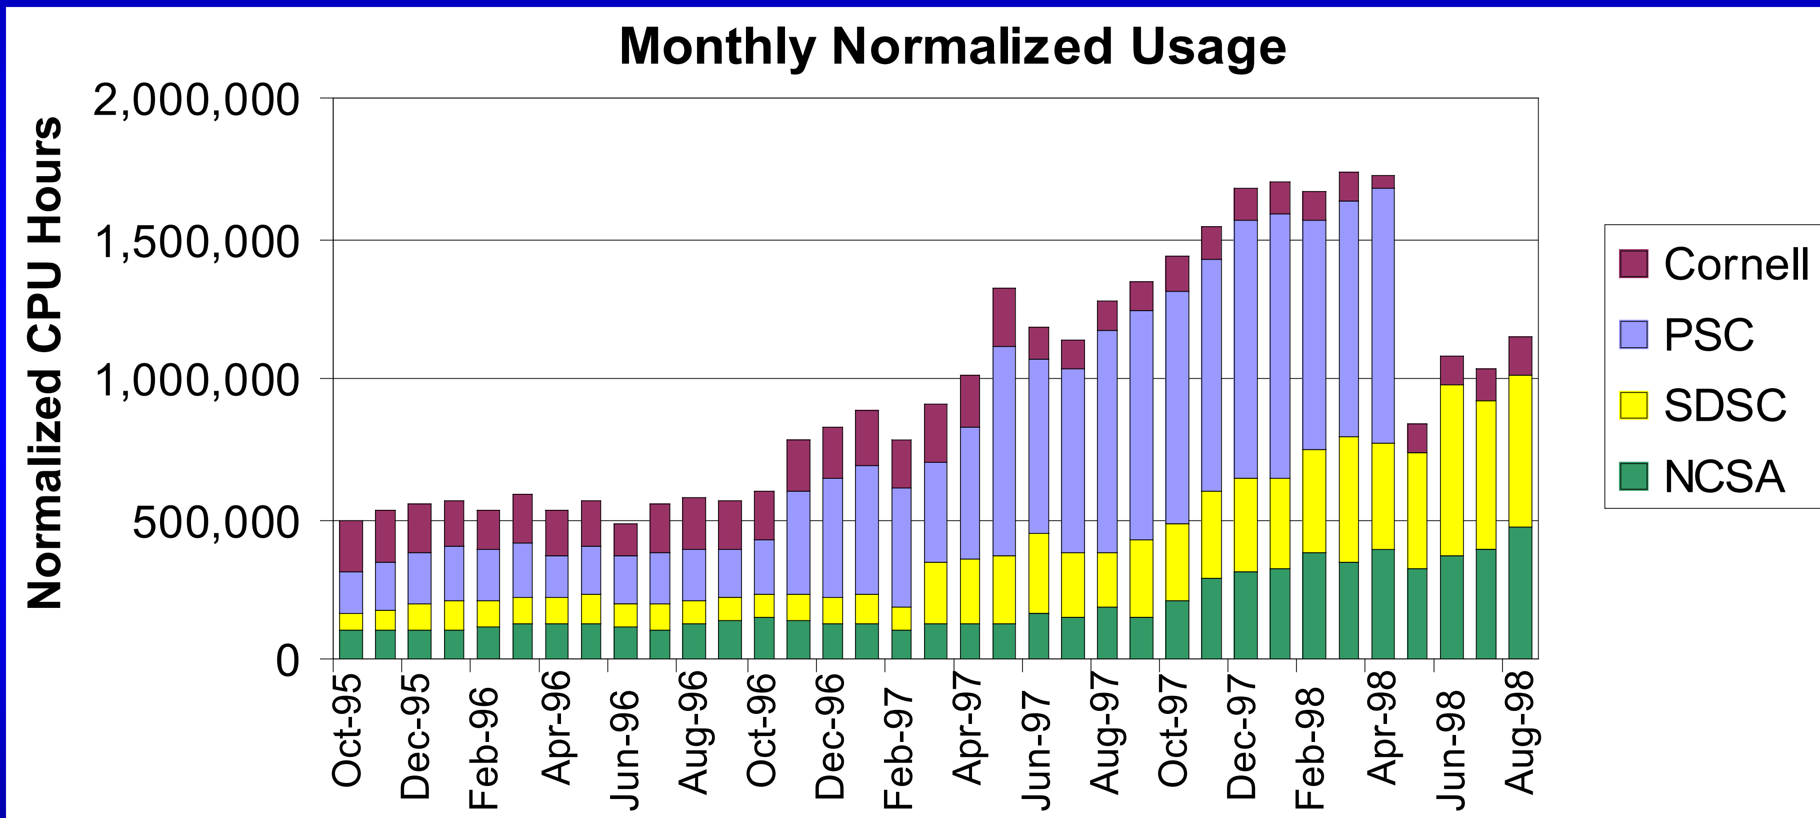

NSF PACI Centers

- Objective is to provide an ASCI level of capacity for the Centers (in the aggregate across the centers)
- Revisit the decision to only have 2 centers.