The Federal HPCC Program: Assessment and Challenge

George Cotter, National Security Agency

Outline

- The Federal High Performance Computing and Communications Program - A Brief Overview
- NSA's Needs; Research and Development Agenda
- HPCC Workshops; Technical Assessments
- Future for the High End User

Federal HPCC Program Goals

- Support to Agency Missions
- Support to HPCC Research Community
- Coordination of R & D Agendas
- Emphasis on Pre-competitive technologies
- Technology Undercarriage to NII
- Demonstration Applications - e.g. Grand Challenges
- Stimulate Academia and Industry to World Leadership

The Federal HPCC Program

COMPONENTS	ACTIVITIES
High Performance Computing Systems	R&D of Scalable Parallel & Distributed Computing Systems, Design Tools, Component Technology Procurement of Advanced Prototypes
Advanced Software Technology and Algorithms	R&D of Operating Systems, Compilers, Languages, Libraries, Data Mgmt, Tools, Algorithms, "Grand Challenge" Applications
National Research and Education Network (NREN)	Extend U.S. Leadership in Computer Communications, Extend Connectivity of, R&D Community to High Performance Computing, Nationwide Prototypes for New Communications Technologies
Basic Research and Human Resources	Grants, Fellowships, Education Programs, Pilot Projects for K-12 & Lifelong Learning
Information Infrastructure Tech. & Applications	Extend HPCC Technologies for National Challenges, Integrate Technologies for Services, Software, and Interfaces

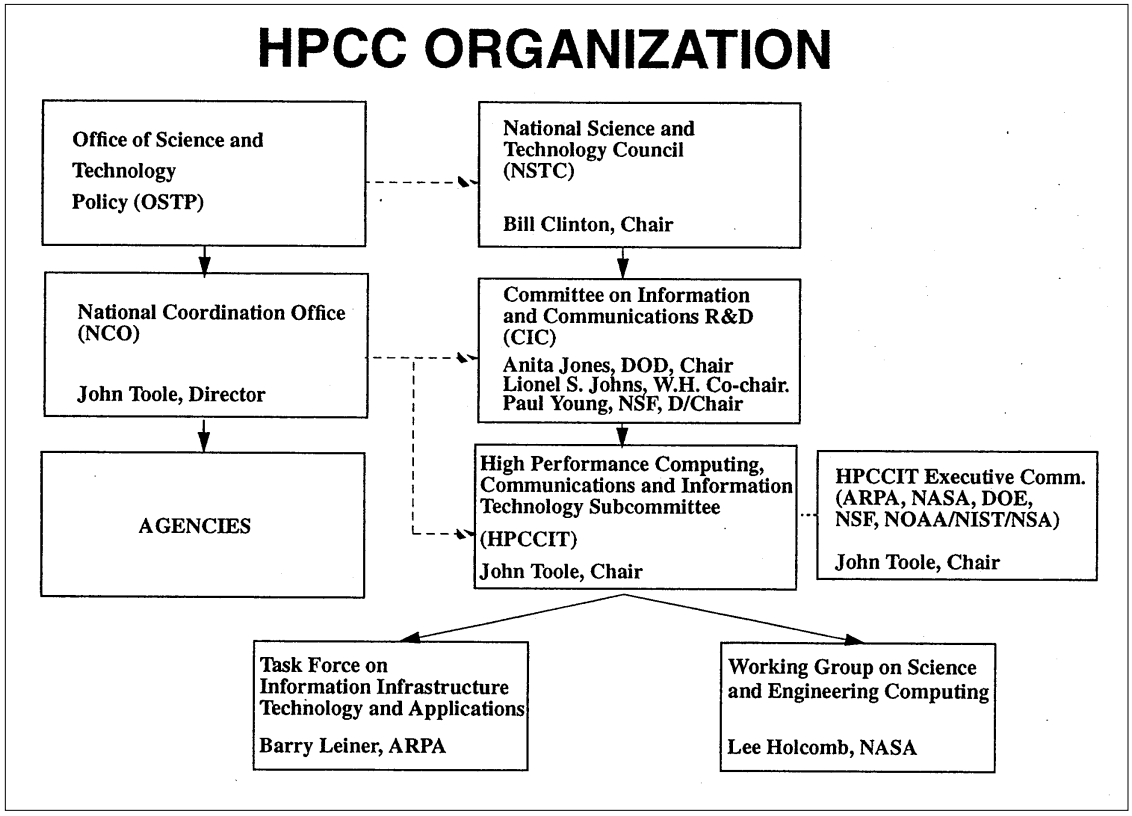
HPCC ORGANIZATION
Office of Science and Technology Policy (OSTP)
National Science and Technology Council (NSTC)
Bill Clinton, Chair
National Coordination Office (NCO)
John Toole, Director
Committee on Information and Communications R&D (CIC)
Anita Jones, DOD, Chair
Lionel S. Johns, W.H. Co-chair.
Paul Young, NSF, D/Chair
AGENCIES
High Performance Computing, Communications and Information Technology Subcommittee
(HPCCIT)
John Toole, Chair
HPCCIT Executive Comm. (ARPA, NASA, DOE, NSF, NOAA/NIST/NSA)
John Toole, Chair
Task Force on Information Infrastructure Technology and Applications
Barry Leiner, ARPA
Working Group on Science and Engineering Computing
Lee Holcomb, NASA

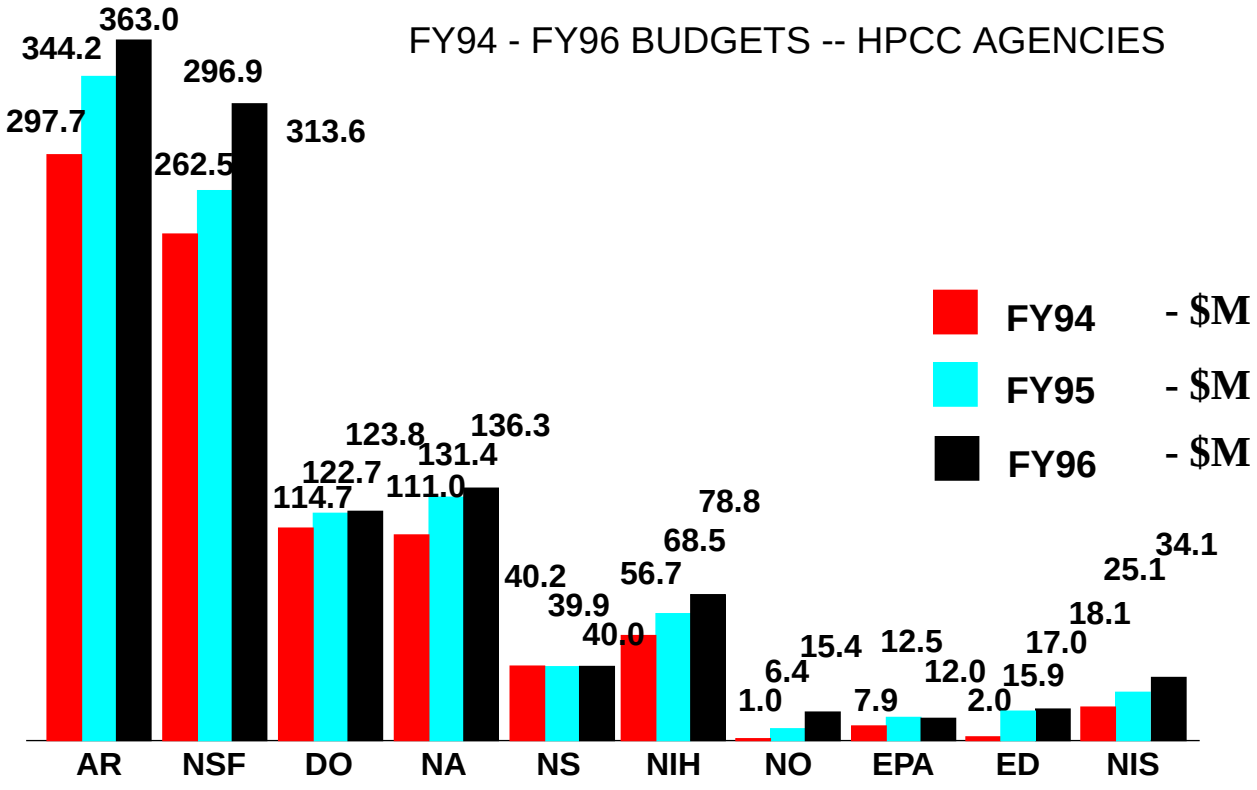
FY94 - FY96 BUDGETS -- HPCC AGENCIES
FY94 - $M
FY95 - $M
FY96 - $M
AR
NSF
DO
NA
NS
NIH
NO
EPA
ED
NIS
297.7
344.2
363.0
262.5
296.9
313.6
114.7
122.7
123.8
111.0
131.4
136.3
40.2
39.9
40.0
56.7
68.5
78.8
1.0
6.4
15.4
7.9
12.5
12.0
2.0
15.9
17.0
18.1
25.1
34.1

HPCC PROGRAM HIGHLIGHTS

HIGH PERFORMANCE COMPUTING

Performance Growth
Embedded Systems
Rapid Prototyping Facility
Microsystems
Networks of Workstations
HPC Architectures

HIGH PERFORMANCE COMMUNICATIONS

Technologies improving Connectivity, Data Rates, Multimedia
Internetworking R&D - Services, Routing & Scaling, Security
Six Gigabit Testbeds - Sustaining 2.4Gbs
Newest Transport Protocols
Contentionless Switches, Host Interfaces
Underpin other R&D - Optical Crossbars, HS Encryptors

MASS STORAGE

10*15, Scalable to 10*18 Bytes
Expanded RAID Architectures
HPSS Consortium/National Storage Lab

ADVANCED SOFTWARE TECHNOLOGIES

Scalable Microkernel Operating Systems
Languages - HPF, HPC++, AC, DBC
PVM, MPI, Ptools

ALGORITHMS, PERFORMANCE MEASUREMENT, BENCHMARKING

RELATED TECHNOLOGIES

NRC/CSTB HPCC REPORT - 1995

FINDINGS:

- ***Information Technology central to our society***
- ***Information Technology advances rapidly***
- ***Retaining IT leadership vital to U.S.***
- ***Federal computing research investment very fruitful***
- ***Continued Federal investment necessary to sustain lead***
- ***HPCC agencies core of government-sponsored HPCC research***

RECOMMENDATIONS:

- ***Continue strong government IT support***
- ***Continue HPCCIT, increasing NII emphasis***
- ***Continue funding research in parallel software/algorithms***
- ***Stop funding development of commercial systems***
- ***Teraflop computer - direction not destination***
- ***Increase communications/networking research***

- ***Start research on very large HP, distributed information systems***
- ***Ensure National Challenges contribute to NII technologies***
- ***Continue NSF centers, funding applications to support HPCC technology***
- ***Ensure Grand Challenges support HPCC technology development***
- ***Strengthen NCO, appoint advisory committee***
- ***Keep projects focused on HPCC objectives***
- ***Mission agency computer procurements only for mission needs***

OUR NEEDS

ULTRA HIGH-END PERFORMANCE

CONTINUOUS HIGH VOLUME THROUGHPUT

BALANCED SYSTEMS

Communications
Computation
Memory

HETEROGENEOUS ARCHITECTURE

Classical Vector Parallel Machines
MIMD-Massively Parallel Processors
Clusters/Symmetric MultiProcessors
Special Purpose Devices
Mass Storage Systems
High Capacity Networks

MATURE SOFTWARE AND TOOLS

Applications Development
Experimental & Production Codes
Very Large User Community

TAXONOMY OF PROBLEM TYPES

HARD

Random Walk

Event Scheduling

Graph Traversal

Ordering

Iteration

Embarrassingly Parallel

EASY

TRENDS - KEY QUESTIONS

Problem Sets vs. Architectures - Mappings

Speed vs. Parallelism - Hedging our Bets

Transition to MPP's - Sooner Or Later?

Hardware Models vs. Programming Models - Facing the Facts

System vs. User Parallelization - A Key Investment Issue

Multitasking vs. Resource Partitioning - Also Facing the Facts

Price vs. Performance - Selling out the User

What Will Market Support - Picking Winners and Losers

Technology Gaps - Spawning R & D

NSA HPCC RESEARCH DIRECTIONS

HARDWARE

PIM
Superconducting Switch
Diamond Substrate Technology
Multi-Chip Module Technology
SPLASH II
H-NET
Gigabit Networking (ATM/SONET)
Evaluation of Architectures
Link/Cell Encryption

SOFTWARE

CM-5 Node C Compiler including Vectors (AC)
Data Parallel, Bit Serial C (dbC)
Large Workstation Clusters (RES)
Cooperative Computing Software
Network File System on a Workstation Cluster (MNFS)
Benchmarks for Parallel Architectures
Secure Microkernel

TOOLS & PERFORMANCE MONITORING

Visualization
Instrumentation of Parallel Architectures

PIM CONCEPT

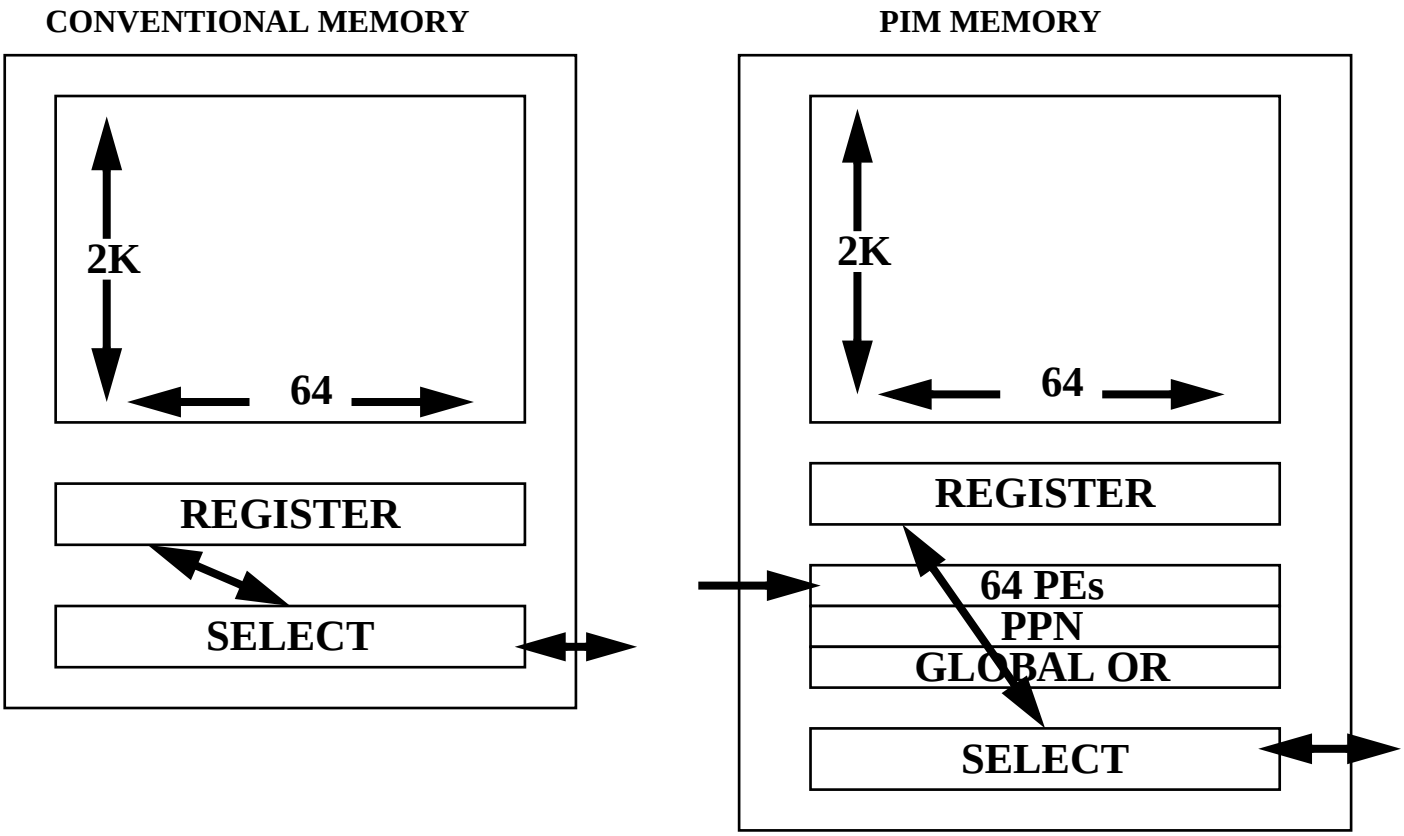

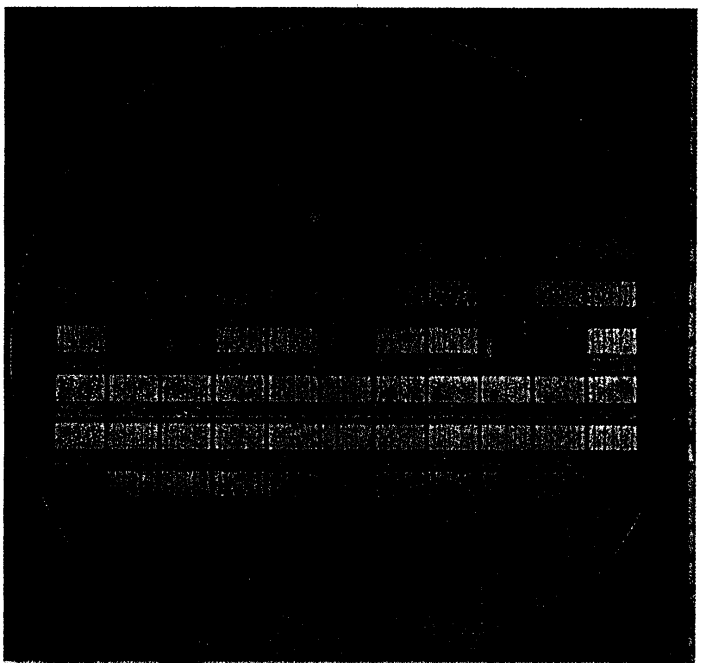

Cray 3/PIM (Processing-in-Memory) Wafer

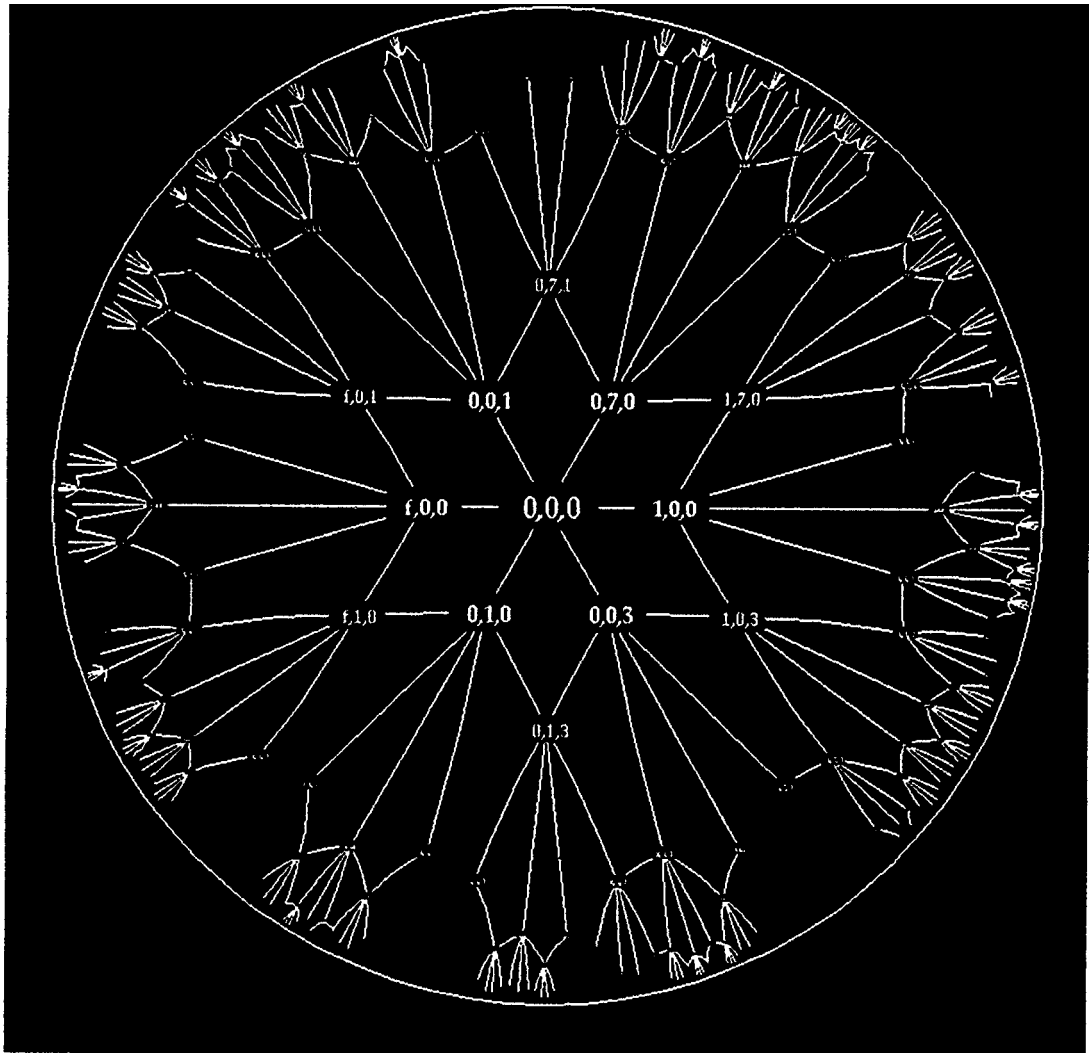

Mapping of three-dimensionsional Toroidal Cray Research T3D network into a two-dimensional workstation screen

UNCLASSIFIED

TIMING THROUGH CROSSBAR FROM PROCESSOR TO MEMORY AND BACK

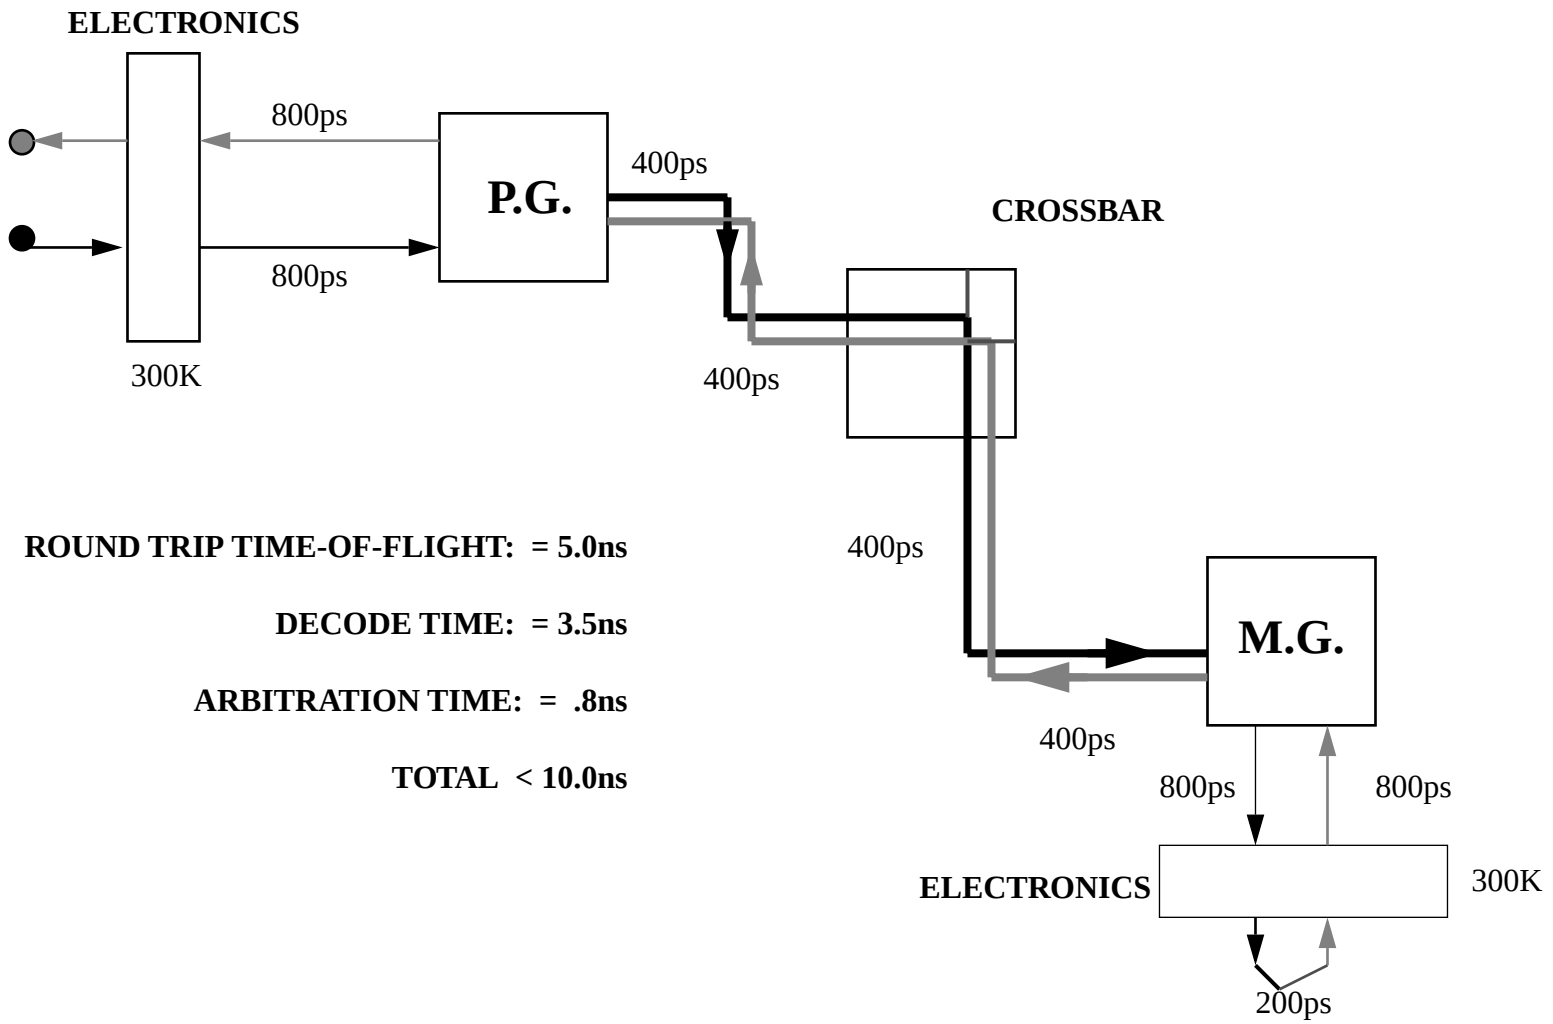

ROUND TRIP TIME-OF-FLIGHT: = 5.0ns

DECODE TIME: = 3.5ns

ARBITRATION TIME: = .8ns

TOTAL < 10.0ns

SOFTWARE TECHNOLOGY FRAMEWORK

Emphasis On Industry Shortcomings

Understanding Architectures; Strengths & Weaknesses
Disconnects between Programming and Hardware Models
Immaturity Of Systems Software Offerings
Tools Built For Less Knowledgeable Users
Uniqueness Of Some Of Our Problems
Support To Our In-house Developments

Examples:

AC-

Node Level Compiler For CM-5
Close Match Between Language & Machine
Used With Any CM-5 Programming Model
10X Improvement Over CM-5 Data Parallel Model
Working With Cray Research For T3D

DBC -

Data Parallel, Bit Serial C
TWIST - Execution Environment
Several Applications - e.g. TERASYS/PIM

Tools, Algorithms, Language Extensions

SOME OBSERVATIONS

Our Computational Challenges Are Boundless

Our Real Limitations Are Intellect, Technology and Money

Heterogeneous, Distributed Computing Is a Fact of Life

Architectural Rules:

First:	*Speed*
Then:	*Balance*
Next:	*Parallel*
Last:	*Massive*

What Isn't Working - Shifting Software Burden Back to User

What Counts:

Implementation - Not Just Invention
Commercialization - We Buy Our Production Systems
Support - Be It Public Domain or Commercial Technology
Growth Potential - Interest Wanes if Migration's Obscure

What's Sad:

Competition Drives Down Profits
Reduced Profit Drives Down R & D
Industry Mortality is too Great
Federal HPCC Program Has Had Little Effect

HPCC WORKSHOP ASSESSMENTS

PURPOSE

Examine Key Areas of High Performance Computing to Help Guide Federal HPCC Program

SPONSORS

Major HPCC Agencies and Collaborators

PARTICIPATION

Government, Industry, Academia Technical Leaders

WORKSHOPS

April 14-16 1992	***Pasadena I - System Software and Tools for High Performance Computing Environments***
May 4-7 1993	***Workshop and Conference on Grand Challenge Applications and Software Technology***
Feb 22-24 1994	***Enabling Technologies for Petaflop Computing***
Jan 10-12 1995	***Pasadena II - System Software and Tools for High Performance Computing Environments***
Aug 14-23 1995	***Petaflops Summer Study on Applications***

HPCC WORKSHOP ASSESSMENTS

PASADENA I - SYSTEM SOFTWARE AND TOOLS FOR HIGH PERFORMANCE COMPUTING ENVIRONMENTS - 14-16 APRIL 1992

KEY FINDINGS

Inadequate Programming Models
Poor Understanding of Resource Balance
Math Libraries Unavailable
Limited User Access to System State
Disparate User Interfaces and Execution Models
Absence Checkpoint/Restate and Exception Handling
Dynamic Resource Management left to Programmer
Programming in Heterogeneous Environment Very Hard
I/O, File Handling Not Keeping Pace

STRATEGY

Unprecedented Cooperation Across HPC Community
Stimulate Complete Programming Environments
Coordinate MultiAgency Research Programs
The Kennedy Model for Advancing Development

HPCC WORKSHOP ASSESSMENTS

PETA(FL)OP I - ENABLING TECHNOLOGIES FOR PETAFLOPS COMPUTING - 22-24 FEB 1994

FOCUS

Applications and Algorithms
Device Technology
Architecture
Software Technology

MAJOR FINDINGS

Cost:	*May be Prohibitive*
Applications:	*Aircraft Industry, Environment, Biology and Medicine, National Security*
Device Technology:	*Silicon Forecast Supports* *Superconductivity Significant* *Optical Systems Unlikely* *Memories on Critical Path*
Architectures:	*I- Shared Memory, Cacheless, Multiprocessor* *II- Massively Parallel Processor* *III- Multigrid PIM (Massive, Fine Grain Parallel)* *System Diameter, Latency Critical Issues*
Software Technology:	*System-level Methodologies Completely Inadequate to Challenge* *Entire New Paradigm May be Essential*

HPCC WORKSHOP ASSESSMENTS

PASADENA II - SYSTEM SOFTWARE AND TOOLS FOR HIGH PERFORMANCE COMPUTING ENVIRONMENTS - JAN 10-12 1995

GOALS

Assess Progress Since Pasadena I

Identify Isues in Development, Deployment and Use of Systems Software and Tools

Develop Recommendations for Resolution, Framework for Continuing Review of Progress

KEY FINDINGS

Systems SW and Tools Lagging Significantly
HPC MPP HW Architectures Dominant Source of Difficulty
Transition from Research to Commercial Products Poor
A Few Bright Spots - MPI, PVM, HPF
HPC Market Will be Driven by Commercial Users
HPC Market Currently Too Small to Attract ISV's
SMP's a New Market Factor - ISV Leverage Potential

PASADENA II - Continued

STRATEGY

Production of Minimum Base Tool Environment
Community Standards Forum - Establish Requirements and Specifications for Tools, Infrastructure Interoperability
Sponsor Detailed Evaluation Studies to Eliminate HW Barriers
Determine Minimum Requirements for HPC HW System
Identify HW Hooks and Mechanisms to Strengthen Systems SW and Tools
Identify Key SW to Exploit ISV Resources
Support Interdisciplinary Investigation of Heterogeneous Computing

IMPLEMENTATION

Create an Oversight Committee for Systems SW and Tools
Create a Standing Committee on HPC SW Tool Set Functionality
Convene a Committee on Infrastructure and Interfaces
Create a "Circle of Excellence" for Sponsorship of Research
Identify Key Applications; Sponsor ISV Porting
Sponsor White Papers; Minimum Requirements for HW Support
Extend Application/System Evaluation/Characterization Studies
Support Rapid Deployment Consortia
Extend NHSE Charter into Broad Community Coordination Role
Create National Distributed Hererogeneous Testbed

HPCC WORKSHOP ASSESSMENTS

PETAFLOPS SUMMER STUDY ON APPLICATIONS - 14-28 AUG 1995 (Preliminary)

REQUIREMENTS

Rich Potential Set of Applications
Some Near Term

APPLICATION CLASSES

Performance Dominated - Memory Scales: N^3/4 Law
Data Intensive - Drives Memory Hierarchy
Balanced Systems - Extend Commercial Development
Exotic Applications

ARCHITECTURAL IMPLICATIONS

Variable Architectures (PETAFLOPS I)
Concurrent Architecture Studies Essential
Agressive Hardware Technology Programs Needed
Integrated Applications-Architecture Studies Must Occur
Detailed Performance Models Required to Verify Architectures

SUGGESTIONS FOR PROGRESS

Complete Applications Case Studies
Detailed Recommendations/Milestones
Follow-up Workshops: Applications/Algoritms, Systems SW, PETAFLOPS II

Table 1: OVERALL ROADMAP TECHNOLOGY CHARACTERISTICS

Year of First DRAM Shipment Minimum Feature (μm)	1995 0.35	1998 0.25	2001 0.18	2004 0.13	2007 0.10	2010 0.07	DRIVER
Memory Bits/Chip (DRAM/Flash) Cost/Bit @volume (millicents)	 64M 0.017	 256M 0.007	 1G 0.003	 4G 0.001	 16G 0.0005	 64G 0.0002	D
Logic (High Volume: Microprocessor) Logic Transistors/cm^2 (packed) Bits/cm^2 (cache SRAM) Cost/Transistor @ volume (millicents)	 4M 2M 1	 7M 6M 0.5	 13M 20M 0.2	 25M 50M 0.1	 50M 100M 0.05	 90M 300M 0.02	L (μP)
Logic (Low Volume: ASIC) Transistors/cm^2 (auto layout) Non-recurring engineering cost/transistor (millicents)	 2M 0.3	 4M 0.1	 7M 0.05	 12M 0.03	 25M 0.02	 40M 0.01	L(A)
Number of Chip I/Os Chip to package (pads) high perf.	 900	 1350	 2000	 2600	 3600	 4800	L, A
Number of Package Pins/Balls Microprocessor/controller ASIC (high performance) Package cost (cents/pin)	 512 750 1.4	 512 1100 1.3	 512 1700 1.1	 512 2200 1.0	 800 3000 0.9	 1024 4000 0.8	 μP A A
Chip Frequency (MHz) On-chip clock, cost-performance On-chip clock, high performance Chip-to-board speed, high performance	 150 300 150	 200 450 200	 300 600 250	 400 800 300	 500 1000 375	 625 1100 475	 μP L
Chip Size (mm^2) DRAM Microprocessor ASIC	 190 250 450	 280 300 660	 420 360 750	 640 430 900	 960 520 1100	 1400 620 1400	 D μP A
Maximum Power High performance with heatsink (W) Logic without heatsink (W/cm^2) Battery (W)	 80 5 2.5	 100 7 2.5	 120 10 3.0	 140 10 3.5	 160 10 4.0	 180 10 4.5	 μP A L

From THE NATIONAL TECHNOLOGY ROADMAP FOR SEMICONDUCTORS, 1994--Semiconductor Industry Association

UNCLASSIFIED

PROJECTIONS FOR THE TERAFLOP COMPUTER IN THE YEAR 2000

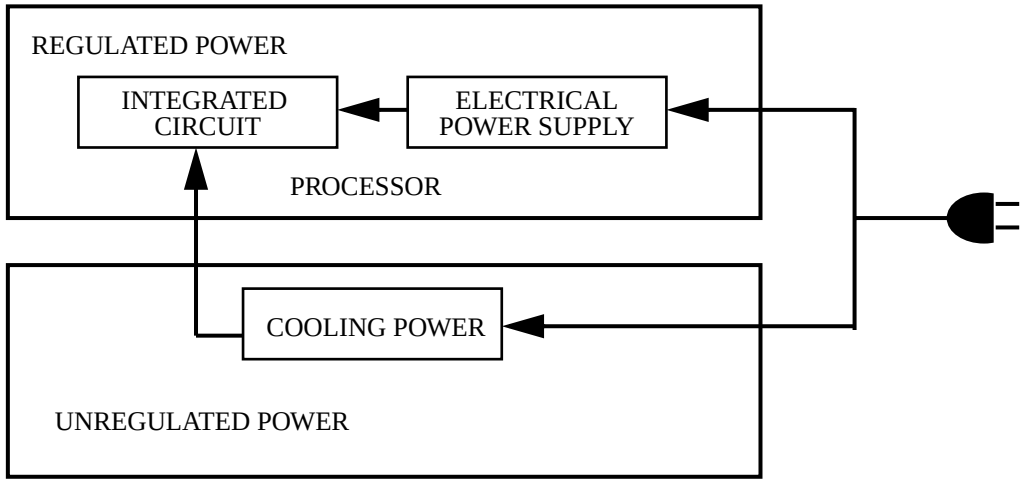

	REGULATED POWER PER UNIPROCESSOR (KW)	NUMBER OF PROCESSORS (#)	"WALL PLUG" POWER (KW)	POWER DENSITY PER UNIPROCESSOR (WATTS/INCH^3)
ECL	1.5	1000	3600	830
GaAs	0.7	500	850	200
JJ	0.001	80 (2.5 Ghz)	200	0.3
JJ	0.001	20 (10 Ghz	200	0.3

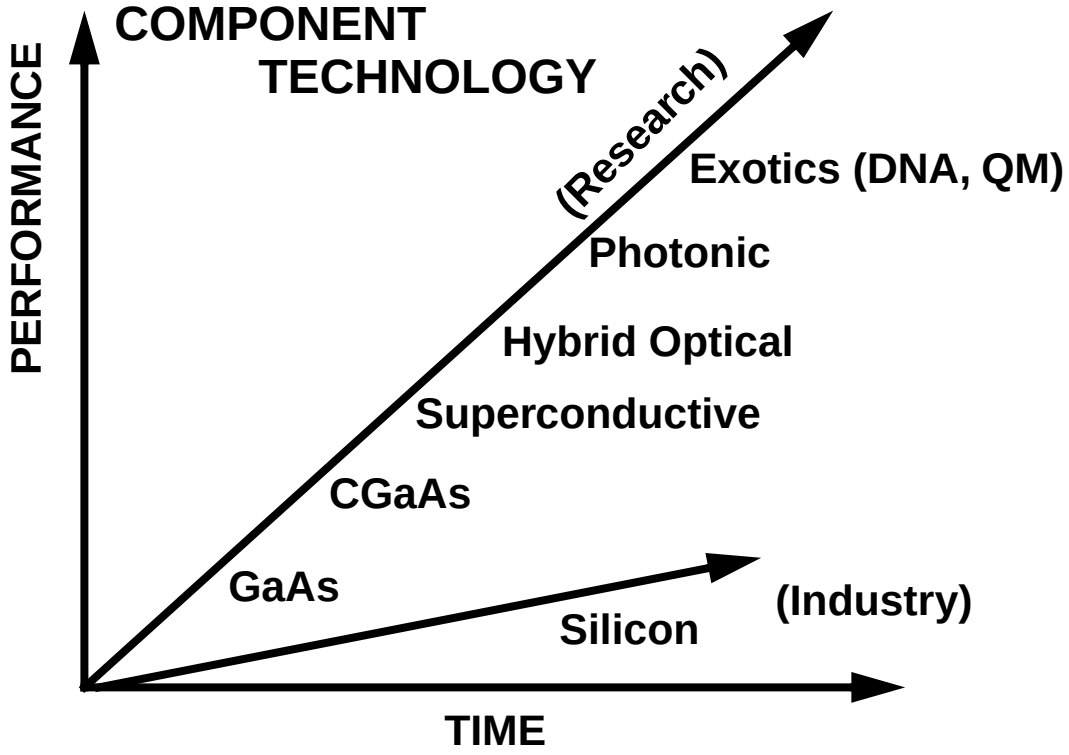
COMPONENT
TECHNOLOGY
PERFORMANCE
(Research)
Exotics (DNA, QM)
Photonic
Hybrid Optical
Superconductive
CGaAs
GaAs
Silicon
(Industry)
TIME

ONE VIEW OF THE FUTURE FOR HPCC

Keep the Performance Gap from Opening Wider

Establish Government-Industry-Academic Consortia

- *Establish Research Program Narrowly Focused-Systems SW*
  - *Language Issues*
  - *Compiler Structures*
  - *Operating Systems (Microkernels to Autoscheduling)*
  - *All Critical, Generic Tool Issues*
  - *User Requirements and Functionality*
  - *Peer-level Selection for Pre-competitive Development*
  - *Conduct Systems-Significant Applications Case Studies*
  - *Incorporate NHSE into Consortia*
- *Create Heterogeneous Computing Testbed*
- *Leverage Systems Vendor and ISV Industries*
- *Objective: Honest Teraflop Performance for General User*

Adopt Petaflop as HPCC Long-range Focus

- *Support Essential Material Science and Device Technology*
  - *Silicon*
  - *Superconductors*
  - *Promising Optical Developments*
  - *Memories a Special Case*
  - *Exotics*
- *Support Architectural Developments (e.g. Petaflop I Models)*
- *Extend Systems Software Consortia Charter*
  - *Research in New Paradigms*
  - *Coordination Applications, Systems Software, Architectural Communities*

OUR CRI EXPERIENCE

- *An Original Client*
- *Acquired Every Machine Developed*
- *Enjoy Broad Technical Collaboration*
  - *C90, T3D, TRITON, T3E Currently*
- *Often Represent CRI Interests Within Government*
- *Seldom at Odds (e.g. Export Control Policy)*
- *Future of CRI Important to the Nation*
  - *National Security*
  - *Competitiveness*
  - *Technology Leadership*