From Megaflop to Petaflop and Beyond

3rd Annual Blue Waters Seminar
Sunriver Resort, Oregon
May 10-13, 2015

Arden L. Bement, Jr.

NCSA

Part A
A Reprise of NSF's Investments in HPC

Peter Lax Report

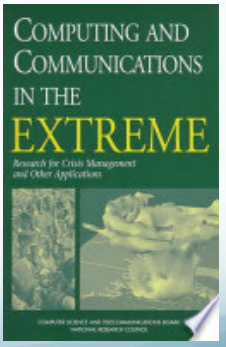

The Lax Report

- Increased access to HPC for S&E
- Increased research in computational mathematics, software and algorithms
- Training of personnel in HPC
- R&D for the implementation of new supercomputing systems

Cray X-MP

The Initial NSF Program

- Cornell Theory Center (CTC)
- The National Center for Supercomputing Applications (NCSA) at UIUC
- The Pittsburgh Supercomputing Center (PSC) CMU, U of Pitt, and Westinghouse
- San Diego Supercomputing Center (SDSC)
- John von Neumann Center (JVNC) Princeton, NJ

Communication Networks

- NSFnet from 1985 to 1995
- vBNS from 1995 to 2001
- Teragrid from 2001 to 2011
- eXtreme Digital Grid 2011-

Study and Planning Activities

- Blue Ribbon Panel (“From Desktop to Teraflop”) 1993
- NRC Committee on High Performance Computing and Communication 1994
- Task Force on the Future of the NSF Supercomputing Centers Program (Hayes Report) 1995

Common Recommendations

- Centers should be judged on scientific impacts
- Centers should be broadly based
- Should be close program and allocation links
- Centers should build regional alliances
- Staff should conduct center-aligned research
- NSF should provide a pyramid of computing capabilities

Hayes Report

- Major Issues
  - MPCs versus vector machines
  - Access to Teraflop capacity
  - Partnership models among centers
- Recommendations
  - More than one Apex partnership centers
  - Apex machines should be MPCs
  - NSF to provide leadership to acquire teraflop capacity at mission agencies

PACI Program

- National Computational Science Alliance led by NCSA
- National Partnership for Advanced Computational Infrastructure (NPACI) led by SDSC
- Four teams driving:
  - Technology developments
  - Building usable tools & infrastructure
  - Distributing computing resources nationally
  - Promoting the use of computational technologies

TeraGrid

- Announced by NSF in 2001
- Involved eleven partner sites
- Coordinated by the Grid Infrastructure Group (GIG) at U. of Chicago
- Provided access to:
  - High-performance computers
  - Data – both online and stored
  - Visualization facilities
  - Other instrumentation

Office of Cyberinfrastructure

- Established in 2005
- Directors:
  - Dr. Deborah Crawford: 2005-2007 coordinated the planning document: *"Cyberinfrastructure Vision for 21st Century"*
  - Dr. Dan Atkins: 2007-2008: led the 2003 blue-ribbon advisory panel report *"Revolutionizing Science and Engineering Through Cyberinfrastructure"*
  - Dr. Ed Seidel: 2008-2009

OCI Track 1 Facilities

- Cray "Blue Waters" HPC at NCSA (13.2 PFLOPS and 415 PBYTES total memory)
- Cray "Kraken" at NICS (1,17 PFLOPS and 147 TBYTES total memory)
- Intel/Dell/Mellanox "Stampede" at TACC (9.6 PFLOPS and 205 TBYTES total memory

NCSA 'Blue Waters' Supercomputer

eXtreme Digital Grid

- NSF establishes technology insertion service team (TIS) with NSCA, TACC, PSC, and NICS.
- New service called eXtreme S&E Discovery Environment (EXSEDE).
- Three new clusters funded:
  - XSEDE Cluster "Gordon" at SDSC (4.8 TBYTES of flash memory)
  - XSEDE Cluster "Backlight" at PSC (16 TBYTES of shared memory)
  - XSEDE Terascale, Lustre-based clusters (2) at CAC

Part B
Looking to the Future

Synaptic Computing
HPC and Data Analysis
Cloud Computing and IOT

Synaptic Computing

IBM's 'True North' Chip

IBM's ''True North Chip''

- ''Brain-inspired'' neural network architecture with 1 million individually programmable neurons and 256 million individually programmable synapsis.
- Designed based on human-scale brain simulation and
- A new progamming paradigm based on a ''corelet'' spiking neuron model

Human Brain Functions

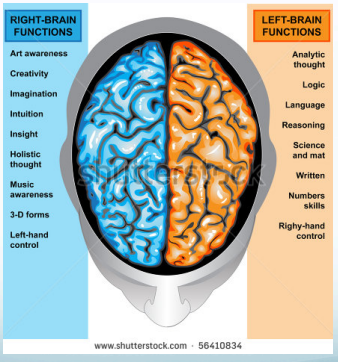

Hybrid Brain-Inspired Computer

- Provides links between Van Neumann architecture (''left brain'') and neural net architecture (''right brain'')
- Opens prospects for forming natural bridges between the human brain and the machine hybrid brain to explore the nebular regions beyond the frontiers of science.

Potential Applications

- Assistive learning, risk analysis, decision making
- Anticipatory controllers
- Bias filter for social research (RCT, DOE)
- Self-aware robots
- Adaptive sensor networks
- Self-operating systems (manufacturing, transportation, public safety, military)
- Smart phone medical diagnostics

HPC and Data Analysis

Solving Grand Challenge Problems

Education and Workforce Development

Grand Challenge Problems

- Usually coupled
- "Nasty" problems (often have more than one solution because of coupled variables)
- Example: coupled energy, environment, economy sustainability
- Problem is bounded by meeting demands for electric power without adding to CO_2 burden or increasing the cost of energy.

NOAA/ U. of Colorado Model

- Requirements:
  - Contiguous U.S. Energy demands met at all times
  - Electrical generation by renewables and natural gas
  - Electrical transmission by high-voltage DC on a national Grid
  - High-temporal and spatial resolution weather data
  - Model year 2030 with 15% increased energy demand
  - New wind generation principally from Texas to Montana/North Dakota Corridor
  - New solar PV generation principally in Souhwest
  - Varied cases of natural gas and renewable electric power ciosts analyzed

Model Results

- Model run on LLL Sequoua 16 PTFLOP Supercomputer
- CO_2 emissions from electric power generation can be reduced by 80% compare with 1990 levels.
- Costs of electricity generation are not substantially increased
- Results can be obtained without electric storage

Education and Workforce Development

Hub Zero Tools:

- Analytical, modeling, and simulation tools
- On-line lectures, courseware, seminars, and practice quizzes
- Information sharing and joint publishing tools
- Community-compiled databases
- Topical articles, publication lists, answers to FAQs

Purdue Hub Group

- Operate 27 hubs, including such communities as:
  - nanotechnology,
  - pharmaceuticals,
  - catalysis,
  - composite materials design
  - Earthquake engineering
  - Cancer science and engineering
  - Biofuels,
  - Climate modeling
  - Water quality
  - Health care
  - High-throughput computing (DiaGrid)

DiaGrid

- Partnership among ten mid-west universities
- Contains 50,000 processors
- Processes approx. 15 million jobs delivering 18 million hours of computing time annually
- Largest regional computing alliance in the US

Cloud Computing and IOT

Cloud Computing and IOT

- Complimentary tools for the control and performance modeling of such complex systems as:
  - Manufacturing
  - Global supply chains
  - Logistic networks
  - Transportation and shipping networks
  - Electric distribution grids
  - Major scientific instruments

South Pole 10-Meter Telescope

Final Note

- Encourage including the following topics in computational S&E courses based on emerging technologies:
  a) Finding solutions to complex, grand-challenge global problems,
  b) Developing on-line, real-time, safe, secure, and reliable controls for complex systems,
  c) Discovering new modalities for the acquisition, assimilation, and application of new knowledge

Thank You