Federal High-End Computing Update

John Grosh (DoD)

Chair, High End Computing
Interagency Working Group

November 17, 2005

High-End Computing Roadmap

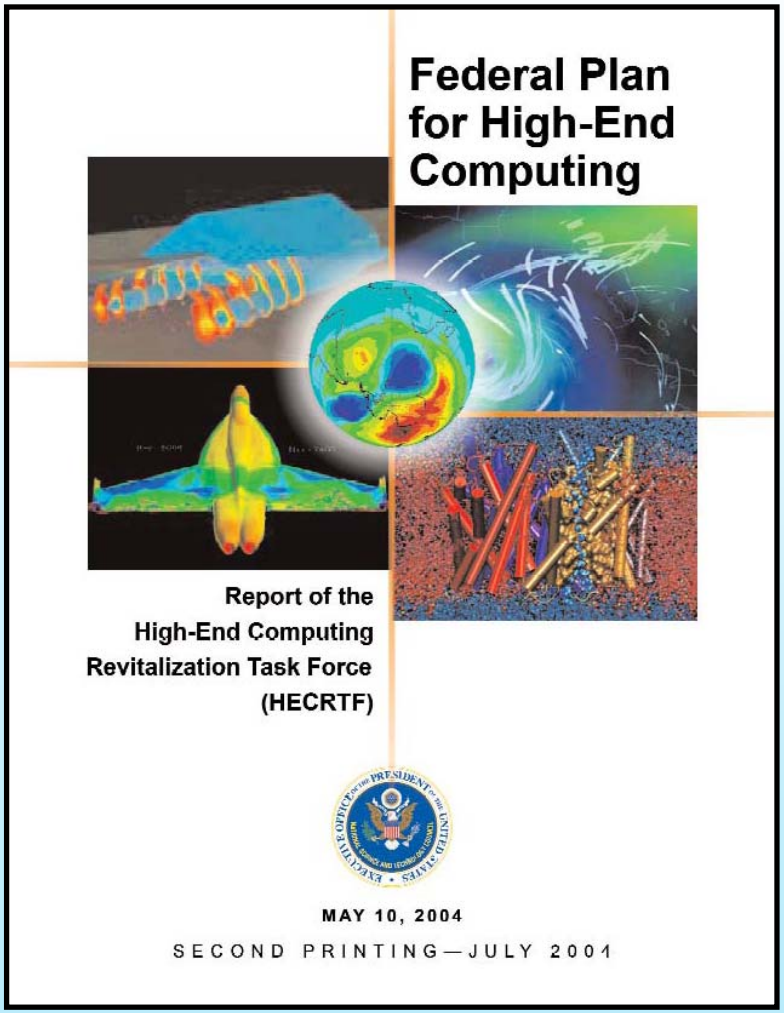

- **Research and Development**
- **Resources**
  - Production Computing
  - Leadership Systems
- **Acquisitions**

Agenda

- **Research and Development**
  - I/O and Storage Workshop (Gary Grider / LANL)
  - HEC University Research Activity and NSF Solicitation in I/O and Storage Research (Almadena Chtchelkanova)
  - DARPA High Productivity Computing System Program / Phase III and Research in Next Generation Language Systems (Robert Graybill)

- **HEC Infrastructure and Resources**
  - DOE Office of Science Leadership Systems / INCITE Program (Barb Helland)
  - NASA Leadership Systems / Solicitation Announcement (Tsengdar Lee)
  - NSF Cyber Infrastructure Plans (José Muñoz)

- **Applications**
  - NIH Roadmap National Centers for Biomedical Computing (Mike Marron)

- **Announcements**
  - DOE NNSA Alliance Program (Robert Meisner)
  - DOE Office of Science SciDAC Program (Fred Johnson)

SC 2005 – HEC BOF

HEC IWG File Systems and I/O Workshop

Gary Grider

Los Alamos National Laboratory

November 17, 2005

Background and Purpose of the Workshop

- **Late in FY04, a multi-agency I/O and File systems research needs document identified a number of needed research focus areas**
- **The HEC IWG decided to make I/O and File Systems a focus area for the coming years and kicked off this effort via the workshop**
- **The Workshop was to provide the first broad input into a multi-year, multi-agency effort to coordinate government spending on research into the HEC IWG focus area, I/O and File Systems**
  - Catalog existing relevant research in this area (high level only) funded by the government, industry, and higher education
  - Look for potential overlaps and gaps
  - List the short and long term top areas that need to be addressed
  - Identified
    - Areas that are receiving enough or too much emphasis
    - Most pressing short term research needs and most strategic long term needs
    - Actions needed to stimulate the next new big ideas in this critical area
    - As soon as possible, a list of short and long term research topics to possibly be included in a near future National Science Foundation (NSF) call for research

Who Was Invited

- **Agencies/Sites**
  - Government HEC Sites
- **Top research institutions in this area**
  - Universities with major emphasis or center of excellence in the file systems and I/O area
- **Select research funding industry leads**
  - ONLY companies that **fund** file systems and I/O research, particularly at the top research institutions, derived from university sponsor lists

The Workshop Identified Categories of Needed Focus

- **Metadata**
  - **Evolution** - Scalability, Extensibility, Archival considerations, ACL's
  - **Revolution** - Scalability, Extensibility, Name Spaces, Archival considerations, Hybrid devices
- **Measurement and Understanding**
  - **Evolution** - Understanding layering contribution, End to end benchmarking and tracing, Visualization
  - **Revolution** - End to end modeling and simulation, VM as tool for large scale simulation
- **QOS**
  - **Evolution** - Determinism with multiple applications and priorities
  - **Revolution** - Adaptive, End to end QOS
- **Security**
  - **Evolution** - Usability, Long term key management, Distributed authentication for file systems, Dealing with security overhead
  - **Revolution** - Novel security as related to file systems and I/O, Novel encryption at rest over time, Key Management, ACL's, End to End encryption API

Categories Continued

- **Next generation I/O architectures**
  - **Evolutionary** - POSIX, Archive considerations, Access aware interfaces, HEC considerations, Small/unaligned I/O, Mixed large and small I/O, Collaborative caching, Impedance matching
  - **Revolutionary** - Redistribution of intelligence and what abstractions we need, Adaptive/reconfigurable stack (application specific perhaps), User space component considerations, File systems that are semantically aware of the data. Novel devices/hybrid devices exploitation
- **File System related communications and protocols**
  - **Evolution** - Exploitation of RDMA/one sided etc., OBSD (transports, security, extensions, applications), NFSv4 (extensions and applications), pNFS (proof of concept, extensions, applications)
  - **Revolution** - Server to server communication
- **Management and RAS**
  - **Evolution -** Reliability and availability at scale end to end and its overhead, Management scaling, Continuous versioning, Power management
  - **Revolution** - Autonomics (adaptive/self healing/predictive), VM as a RAS enabler, Novel devices as enabler
- **Archive (as it relates to I/O and File Systems)**
  - Content addressable, Deep archive on disk, Object archives/parallel archives, Scheduling movement/ILM
- **Assisting research**
  - Testbeds, Clearing houses (providing traces, reliability info, etc.), Support growth of I/O students

Related Materials and Plans

- You will find workshop presentations, a preliminary inter-agency suggested research document, the final report and other referenced materials at **http://www.nitrd.gov/subcommittee/hec/workshop/20050816_storage/**
- The related NSF call for IO and FS research can be found at **http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=13645&org=CCF**
- HEC POSIX I/O API enhancement effort is occurring, see talk at NNSA/ASC booth Thursday afternoon
- Future plans for this effort are currently under discussion within the HEC IWG

SC 2005 – HEC BOF

High-End Computing University Research Activity NSF Solicitation in I/O, File systems and Storage Research

Almadena Chtchelkanova

National Science Foundation

November 17, 2005

NSF HECURA 2004

- FY 2004 NSF/DARPA activity focused on research in
  - Languages
  - Compilers
  - Libraries
- 100 proposals submitted in July 2004
  - 82 projects submitted by 57 US academic institutions and non-profit organizations
    - Includes no-cost national lab and industrial lab collaborators
- Nine projects were awarded
  - Tools and libraries for high-end computing
  - Resource management
  - Reliability of high-end systems

NSF HECURA – 2006 FOCUS

- I/O, file and storage systems design for efficient, high throughput data storage, retrieval and management in the HEC environment.
- hardware and software tools for design, simulation, benchmarking, performance measurement and tuning of file and storage systems.

HECURA – 2005 SCOPE

- File Systems Research
- Future File Systems related protocols
- I/O middleware
- Quality of Service
- Security
- Management, reliability, and availability at scale
- Archives/Backups as extensions to file systems
- Novel storage devices for the IO stack
- I/O Architectures
- Hardware and software tools for design, simulation of I/O, file and storage systems.
- Efficient benchmarking, tracing, performance measurement and tuning tools of I/O, file and storage systems

HECURA Planning Group

- Candy Culhane (NSA) – Chair
- Almadena Chtchelkanova, Jose Munoz (NSF)
- Robert Graybill (DARPA)
- John Grosh (DOD)
- Fred Johnson (DOE/SC)
- Robert Meisner, Thuc Hoang (DOE/NNSA)

SC 2005 – HEC BOF

High Productivity Computing Systems Program Update

Bob Graybill

DARPA

November 17, 2005

HPCS Program Phases I - III

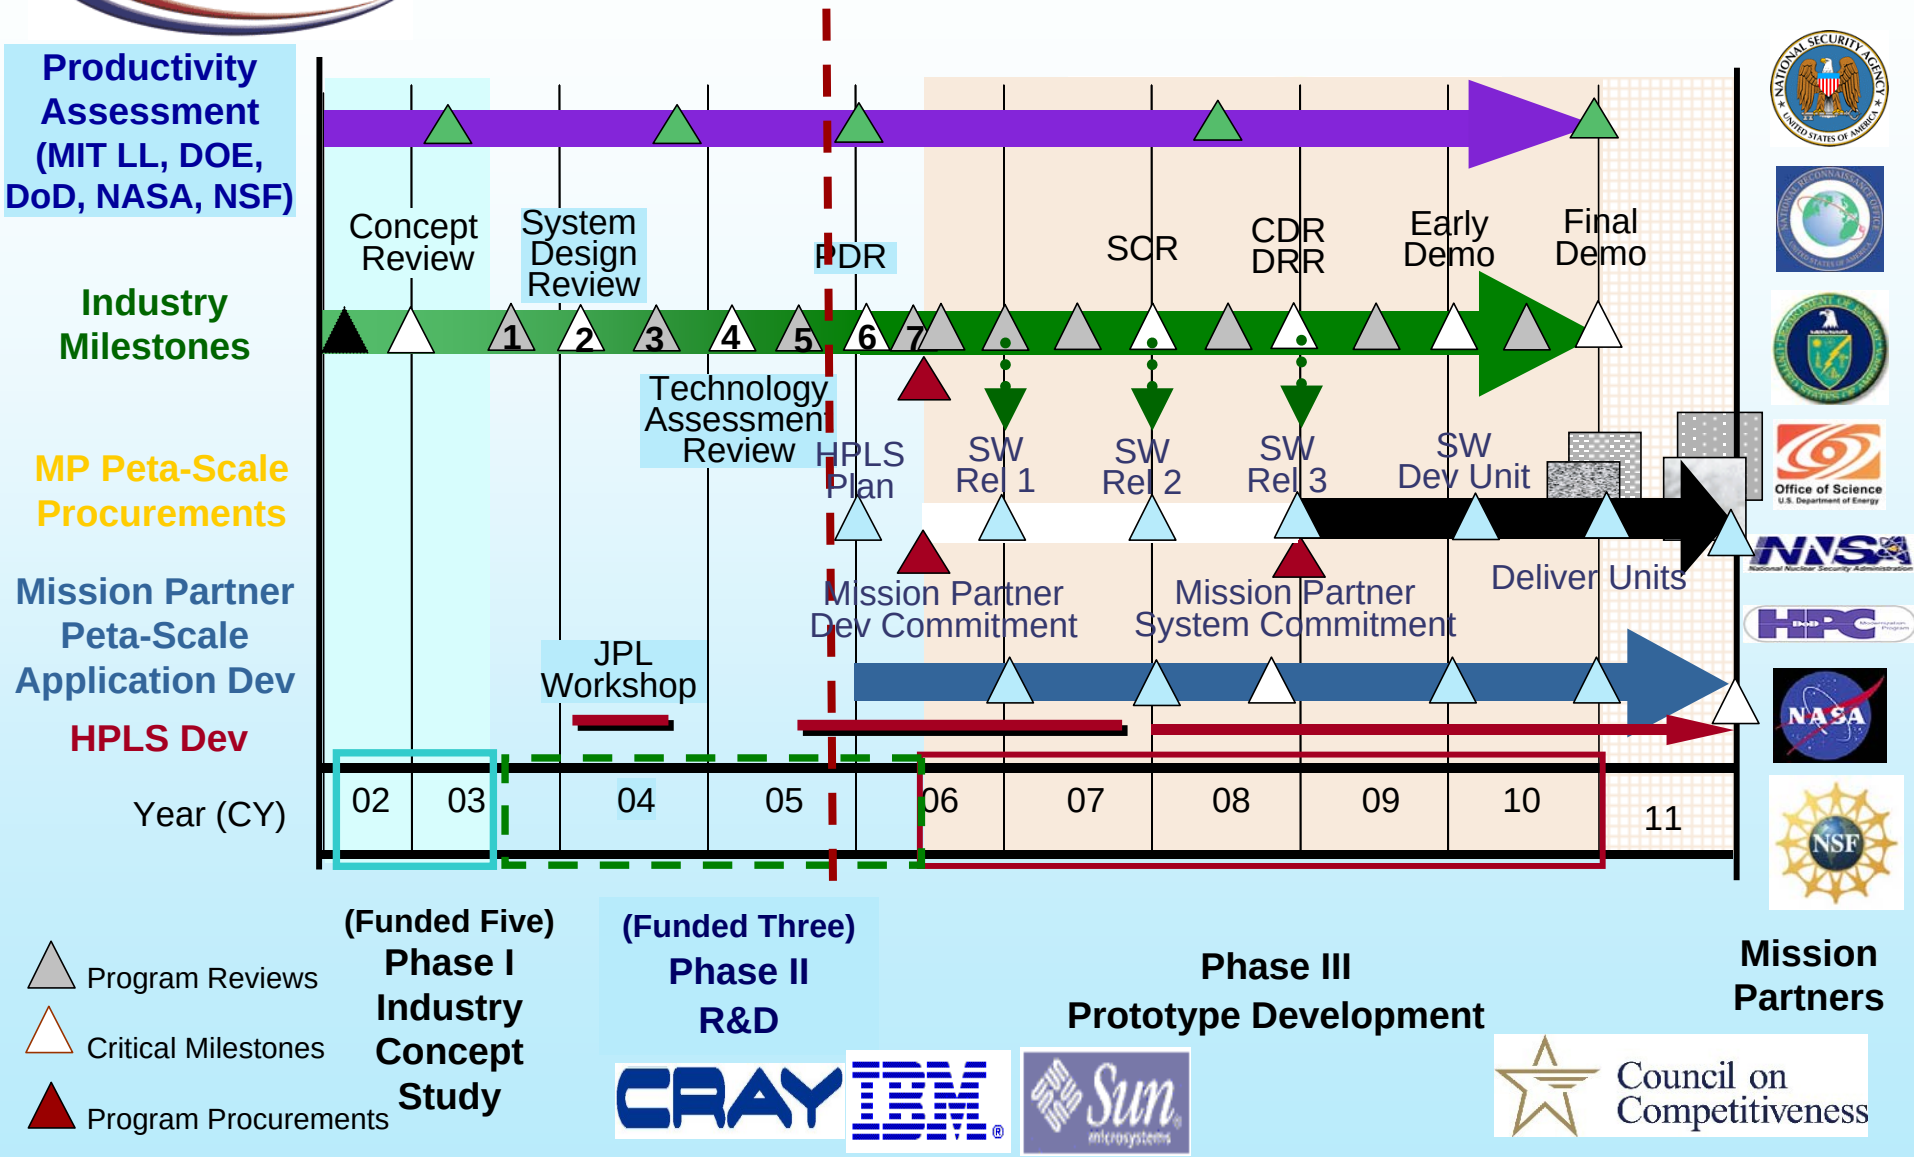

Programming Models Project for Petascale Systems

- An evolutionary approach to building a foundation for an innovative new high productivity language ecosystem
- Team: Argonne, Berkeley, Oak Ridge, Rice, NWU
- Sponsors: DARPA, DOE, and ARDA

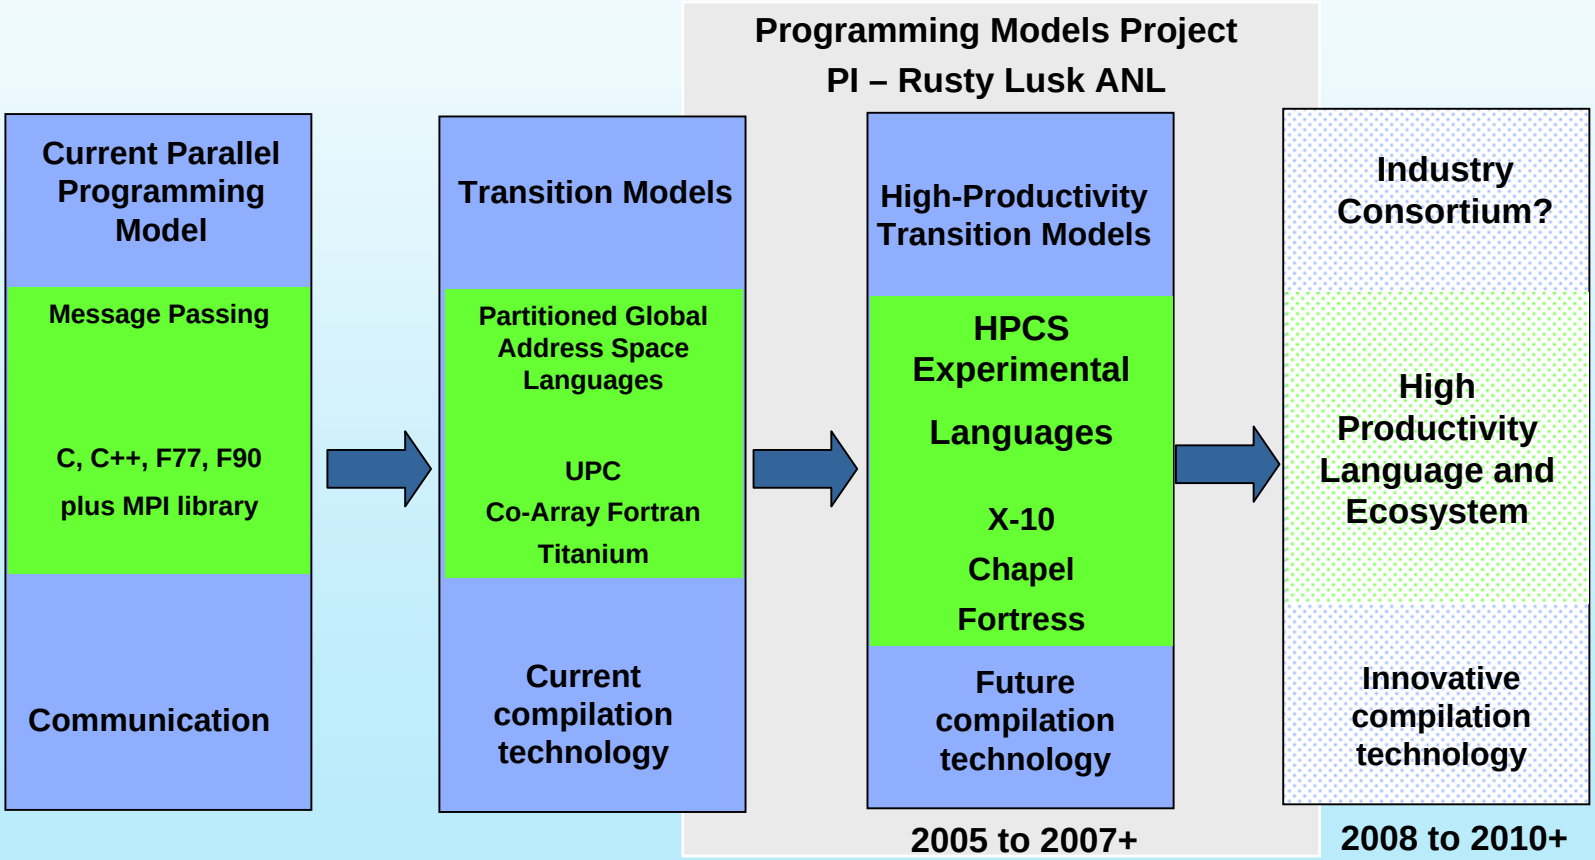

SC 2005 – HEC BOF

Leadership Computing and INCITE

Barbara Helland

Department of Energy/Office of Science

November 17, 2005

Leadership Computing: A Priority at DOE

- In **November, 2003** Secretary Abraham announces 20 Year Science Facility plan: #2 Near Term Priority – UltraScale Scientific Computing Capability
- DOE Office of Science issued to SC laboratories a call for proposals to provide Leadership Class Computing Capability for Science with funding profile of $25M/year for five years in **February, 2004**.
- Four highly competitive proposals received in **April, 2004**

On May 12, 2004 *Secretary Abraham announces the establishment the Leadership Computing Facility at Oak Ridge National Laboratory in Tennessee*

LCF: Current Machines

Cray X1E

18.5 Teraflops
1,024 processors
Custom Vector Processors, Globally Addressable Memory

Used for large projects in Chemistry, Astrophysics, Combustion, Fusion and Accelerator Simulations

Cray XT3

25 Teraflops
5,294 processors

Massively parallel, 3-D Torus Interconnect, AMD Opteron™ processors

Initial LINPACK numbers: ***20.53 Tflops***

Early Test Users: Fusion, Combustion and Astrophysics

- **Next Steps:**
  - **--Cray X1E and XT3 accepted and in production**
- **Expansion is budget dependent**
- **Strong Congressional support**

Expanding the Impact of High-performance computing on U.S. Science

Innovative and Novel Computational Impact on Theory and Experiment- INCITE

- Initiated in 2004, INCITE is open to all scientific researchers and research organizations, including industry.
- Emphasizes computationally intensive research projects of large scale, that can make high-impact scientific advances through large allocation of computer time and data storage. No need for other DOE support.
- For the 2004 and 2005 allocations, 75 proposals were submitted requesting a total of 160 million hours of processing time. Approximately 65% of the proposals submitted were from academic researchers and over 45% of the projects were supported by research agencies other than DOE
- **FY2004 Awards: 4.9 Million processor hours at NERSC awarded to three projects**
  - **Thermonuclear Supernovae: Stellar Explosions in Three Dimensions – Plewa, University of Chicago, 2.7Million processor hours**
  - **Fluid Turbulence and Mixing at High Reynolds Number – Yeung, Georgia Tech, 1.2 Million processor hours**
  - **Quantum Monte Carlo Study of Photoprotection via Carotenoids in Photosynthetic Centers – Lester, UC Berkeley, 1.0 Million processor hours**
- **FY2005 Awards 6.5 Million processor hours at NERSC awarded to three projects December, 2004**
  - **Direct Numerical Simulation of Turbulent Non-premixed Combustion - Fundamental Insights towards Predictive Modeling** – Chen of Sandia National Laboratories in Livermore, California. **2.5 Million processor hours**
  - **Magneto-rotational instability and turbulent angular momentum transport** – Cattaneo, University of Chicago. **2 Million processor hours**
  - **Molecular Dynameomics** – Daggett, University of Washington. **2 Million processor hours**

FY2006 INCITE

- **Expand to include high end computing resources from four DOE National Laboratories:**
  - **5% of the HP MPP system at PNNL**
  - **10% of the IBM Power 3 at LBNL**
  - **10% of Cray Leadership-class computers at ORNL,**
  - **10% of the IBM BlueGene at ANL**
- **Call for Proposals issued May 16, 2005**
- **Award announcement is pending**
- **Look for future announcements at hpc.science.doe.gov**

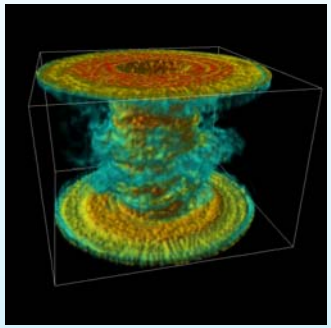

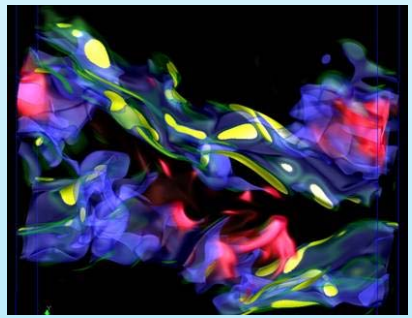

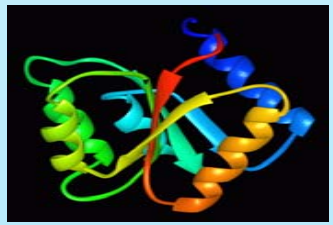

SC 2005 – HEC BOF

NASA's National Leadership Computing System (NLCS)

Tsengdar Lee
Rupak Biswas

National Aeronautics and Space Administration

November 17, 2005

Characteristics of NLCS

- **Portion of NASA's 10,240 processor Columbia Supercomputer system**
  - 20 SGI Altix SSI nodes, each with 512 CPUs and 1TB memory, interconnected via Voltaire® InfiniBand fabric
- **Four of the Altix 3700 Bx2 nodes linked to form a 2048-processor shared memory environment (NLCS)**
  - Intel Itanium2 processors (1.6 GHz clock, 9MB level3 cache)
  - SGI NUMAlink4 interconnect
- **Suited for large-scale applications that can exploit high degree of concurrency and involve substantial inter-processor communication**

NASA's NLCS Call for Proposals

- Released November 2, 2005
- Proposal due January 16, 2006
- Selection expected by March 1, 2006
- Part of NASA's contribution to the 2004 Federal Plan for High-End Computing
- Encouraged by Council on Competitiveness
- Modeling the DOE INCITE program
- Open to National Agencies, academia, and industry
- 4.5 million processor-hours available per year
- 4-8 projects anticipated
- http://www.nas.nasa.gov/Users/Accounts/pi_NLCS.html

SC 2005 – HEC BOF

NSF CyberInfrastructure Plans

José L. Muñoz, Ph.D.

National Science Foundation

November 17, 2005

Office of CyberInfrastructure

- **New Office of CyberInfrastructure (OCI) created in the Office of the Director**
  - Division of Shared CyberInfrastructure (SCI) in CISE no longer exists
- **NSF's plans for CyberInfrastructure may be found in: NSF's Cyberinfrastructure Vision for 21st Century Discovery**
  - Consists of four chapters covering the period (2006-2010):
    - Strategic plan for High Performance Computing
    - Strategic plan for Data, Analysis and Visualization
    - Strategic plan for Collaboratories, Observatories and Virtual Organizations
    - Strategic plan for Education and the Workforce

HPC Acquisition Activities

- **HPC acquisition will be driven by the needs of the S&E community**
- **RFI held for interested Resource Providers and HPC vendors on 9 Sep 2005**
- **First in a series of HPC S&E requirements workshops held 20-21 Sep 2005**
  - Generated Application Benchmark Questionnaire
  - Attended by 77 scientists and engineers

HPC Acquisition

- **05-625 High Performance Computing System Acquisition: Towards a Petascale Computing Environment for Science and Engineering**
  - Due date: 10 Feb 2006
  - Resource Provider/Vendor joint proposal
  - Up to two HPC systems will be acquired
    - Up to $30M for system acquisition
  - Operations and Maintenance costs for RP funded separately
  - Set of benchmarks will be required to be executed and results presented
  - First in a series of acquisitions of this nature
  - Proposals required to address HPC system and O&M
- **There will be a separate announcement for a sustained petascale system acquisition in late FY2006**
- **http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=13649**

SC 2005 – HEC BOF

NIH Roadmap National Centers for Biomedical Computing

Mike Marron

National Institutes of Health

November 17, 2005

Current NIH Activities

- International networks for biomedical data & software sharing: caBIG (NIH/NCI), BIRN (NIH/NCRR)
- National Health Information Network support w/ HHS/ONCHIT
- NIH Roadmap National Centers for Biomedical Computing (NCBCs)
- Cancer Imaging and Computational Centers
- Modeling of Infectious Disease (MIDAS)
- Bioinformatics Resource Centers for Emerging and Re-emerging Infectious Disease
- Proteomics and Protein Structure Initiatives
- P41 Computational Centers – nine exist, on tools & cyberinfrastructure
- National Library of Medicine information & analysis servers
- Many Investigator Research Grants contain supercomputing support
- [more]

Construct an environment for comprehensive mapping and analysis of molecular cellular interactions

- Evidence integration framework to collect and fuse diverse cellular interaction clues based on statistical relevance
- Comprehensive set of physics- and knowledge-based methodologies to fill framework
- Methodologies and filters, anchored in formal domain ontologies, to associate specific interactions to an organisms, tissue, molecular, and cellular context

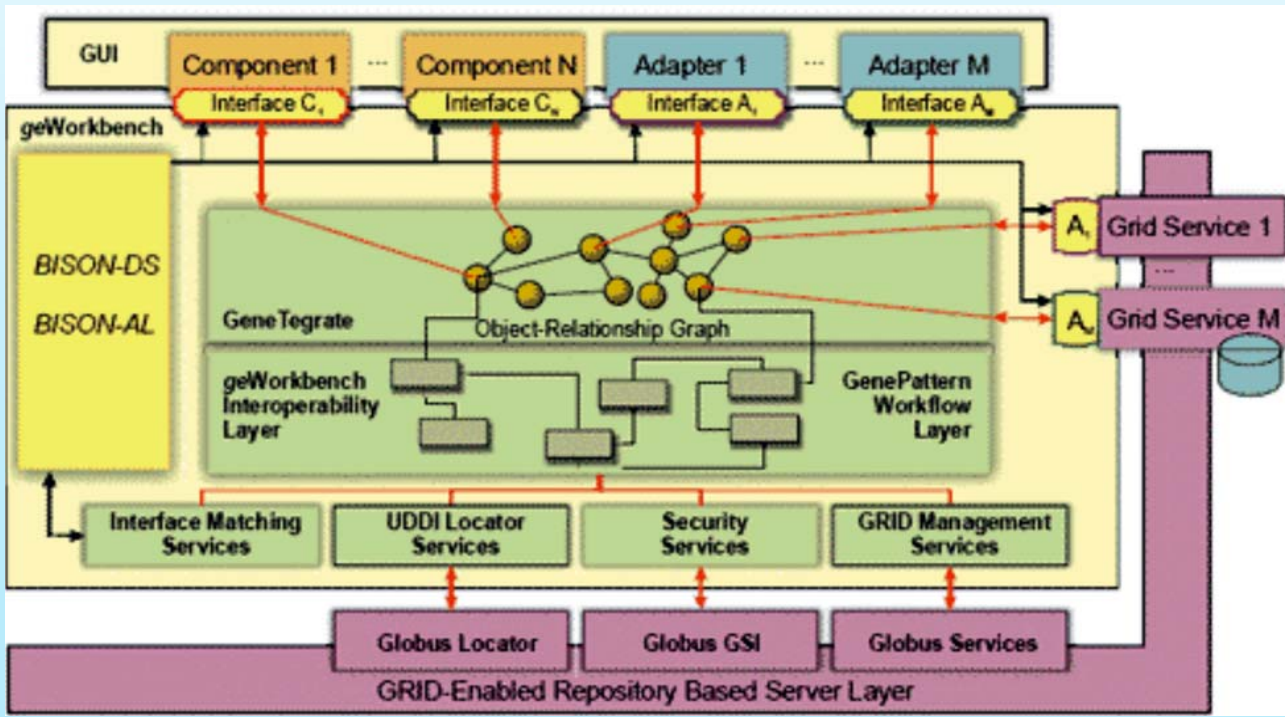

Physics-based Simulation of Biological Structures (SIMBIOS)

PI: Russ Altman, MD, PhD
PI Institution: Stanford University

This Center will develop, disseminate, and support a simulation tool kit (SimTK) that will enable biomedical scientists to develop and share accurate models and simulations of biological structures from atoms to organisms. SimTK will be an open-source, extensible, object-oriented framework for manipulating data, models, and simulations. The software will include advanced capabilities for modeling the geometry and physics of biological systems, generating the governing differential equations of these systems, integrating the equations to simulate the system dynamics, and interpreting the simulation results through comparison with experimental data.

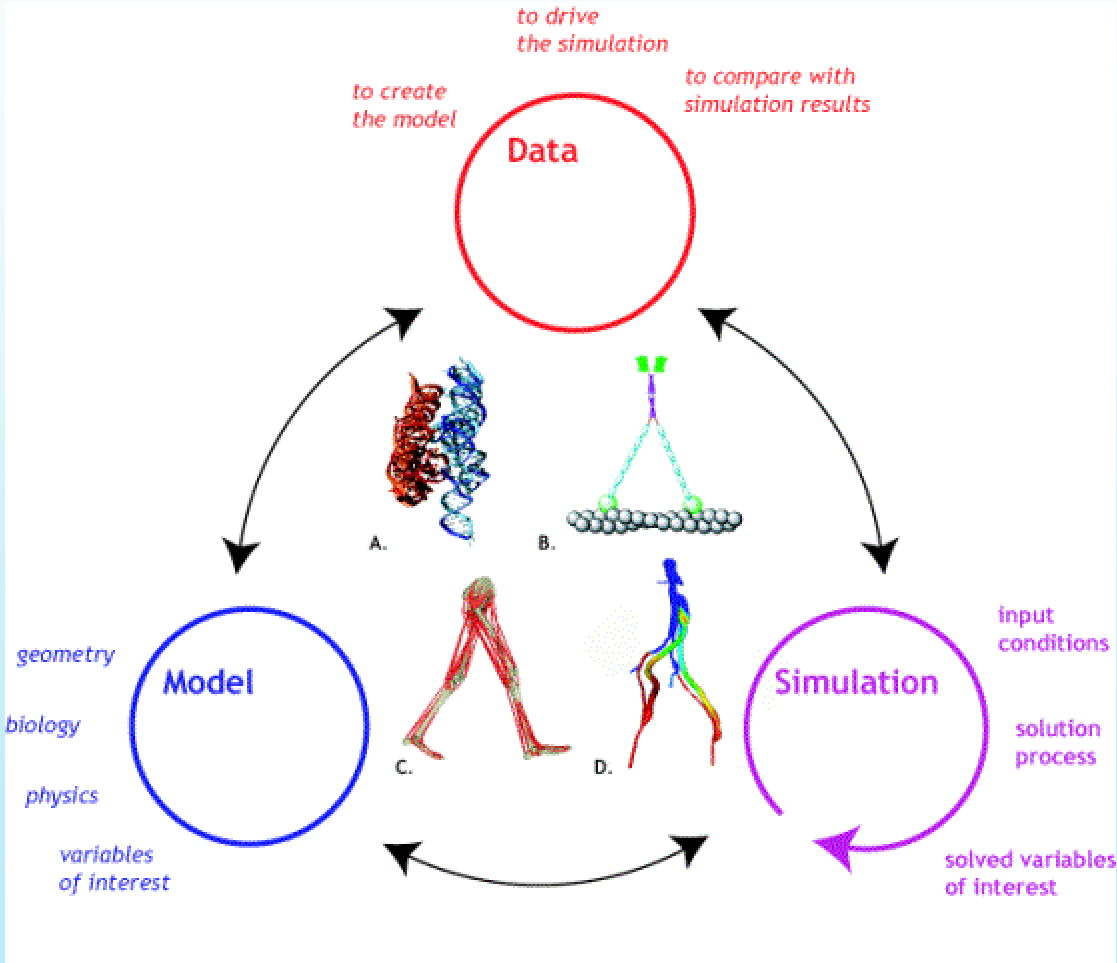

NITRD

MAGNet MULTISCALE ANALYSIS OF GENOMIC AND CELLULAR NETWORKS

http://magnet.c2b2.columbia.edu/

i2b2
Informatics for Integrating Biology & the Bedside

http://www.i2b2.org

NCIBI
NATIONAL CENTER FOR INTEGRATIVE BIOMEDICAL INFORMATICS

http://www.ncibi.org/

CENTER FOR COMPUTATIONAL BIOLOGY

http://www.loni.ucla.edu/CCB/

cBi The National Center for Biomedical Ontology

http://bioontology.org/

National Alliance for Medical Image Computing

http://www.na-mic.org

simbios Stanford Center for Biomedical Computation

Simbios - Physics-Based *Sim*ulation of *Bio*logical Structures

http://simbios.stanford.edu/

SC 2005 – HEC BOF

DOE NNSA Predictive Science Academic Alliance Program (PSAAP)

Bob Meisner

Department of Energy (DOE)
National Nuclear Security Administration (NNSA)

November 17, 2005

Transitioning from the Current to the New Academic Alliance Program

CURRENT ASAP PROGRAM:

- Demonstrate value of validated, large-scale computational science on problems of national importance.
- Develop tightly coupled, multi-physics, 3-D simulations on ASC systems with thousands of processors and initiate verification and validation.

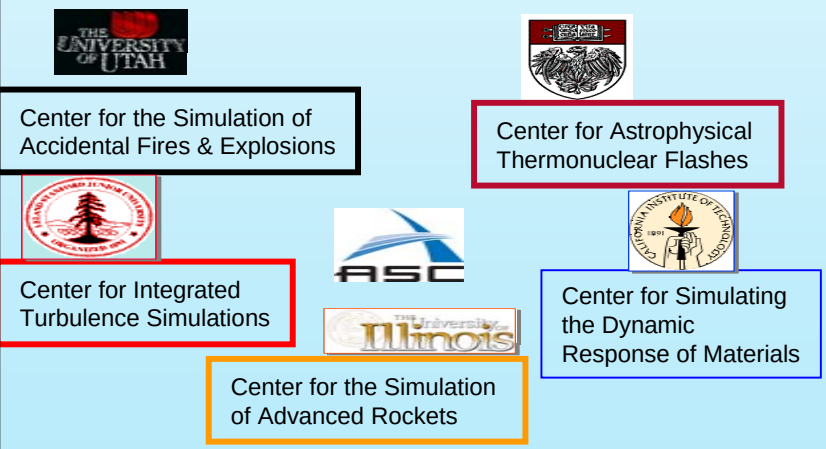

NEW PSAAP PROGRAM:

- Establish validated, large-scale, multidisciplinary, simulation-based "Science-of-Prediction" as a major academic and applied research focus area.
- Predictions derived from such simulations must be strongly connected with experimental and/or observational data.

Proposal Requirements for the New Academic Alliance Program

- **It's mandatory that the proposals address the following two topics:**
  - Predictability in science & engineering.
  - Verification & validation strategies for large scale simulations, including quantification of uncertainty and numerical convergence.
- **The proposed applications and associated sub-disciplines are required to have a strong direct connection to NNSA interests.**
- **Visit the ASC Web site in April 2006 for PSAAP announcements:**

http://www.sandia.gov/NNSA/ASC/

SC 2005 – HEC BOF

Scientific Discovery through Advanced Computing (SciDAC): Status and Update

Fred Johnson

Department of Energy/Office of Science

November 17, 2005

SciDAC FAQ

- **Started in 2001**
- **Goal:**
  - Advance scientific discovery through the use of terascale modeling, simulation and data analysis
- **Strategy**
  - Application teams, laboratory/academic partnerships
  - Mathematical tools
  - Software infrastructure
  - Support for “Big Data” and distributed collaboration
- **Level of effort**
  - $60M/year for five years
  - Applications spanning Basic Energy Sciences, Biological and Environmental Research, High Energy Physics, Nuclear Physics, Fusion Energy Sciences

Status

- **Awards**
  - http://www.osti.gov/scidac/
- **Scientific Progress**
  - 2005 SciDAC meeting, June 26-30, San Francisco
    - http://www.csm.ornl.gov/workshops/SciDAC2005/SD-home.html
  - Institute of Physics online conference proceedings
    - http://www.iop.org/EJ/toc/1742-6596/16/1

Next Steps

- **Recompetition: SciDAC-II**
- **Announcement issued in 4 to 6 weeks**
- **Interagency participation**
  - NNSA, active participant
  - NSF, active discussion
- **Watch Grant Solicitation Notices at:**
  - http://www.sc.doe.gov/grants/grants.html