Update on Scientific Challenges Workshop Series

ASCAC meeting

Paul Messina

March 4, 2009

Outline

- Purpose of workshops
- General structure, process, outputs
  - Follow Basic Research Needs (BRN) workshops model for community-led workshops
  - Use of quad chart templates
  - White papers
- List of all “Scientific Challenges” workshops
- Examples of findings from the three workshops that have taken place
- Related and follow-on workshops
- Credits where credit is due

Purpose of workshops

- To identify grand challenge scientific problems in [research area] that can exploit computing at extreme scales to bring about dramatic progress toward their resolution.
- The goals of the workshops are to
  - identify grand challenge scientific problems […] that could be aided by computing at the extreme scale over the next decade;
  - identify associated specifics of how and why new high performance computing capability will address issues at the frontiers of […]; and
  - provide a forum for exchange of ideas among application scientists, computer scientists, and applied mathematicians to maximize the use of extreme scale computing for enabling advances and discovery in […].

Process used

- Workshop chair(s) work with relevant DOE program offices and colleagues to identify key areas to cover
- Panels defined (4 – 6), panel chairs recruited
- White papers for each panel drafted and posted in advance of workshop
- Priority Research Directions (PRDs) identified by each panel
- Panels populated by domain science experts as well as mathematicians and computer scientists
- Observers from agencies and math and CS community invited to each workshop
- Panels report in plenary sessions in middle and at end of workshop

Priority Research Direction (use one slide for each)

Scientific and computational challenges

Brief overview of the underlying scientific and computational challenges

Summary of research direction

What will you do to address the challenges?

Potential scientific impact

What new scientific discoveries will result?

What new methods and techniques will be developed?

Potential impact on SCIENCE DOMAIN

How will this impact key open issues in SCIENCE DOMAIN?

What's the timescale in which that impact may be felt?

Challenges in Climate Change Science and the Role of Computing at the Extreme Scale

November 6-7, 2008, DC

- Chair: Warren Washington
- Panel topics
  - Model Development and Integrated Assessment
  - Algorithms and Computational Environment
  - Data, Visualization and Productivity
  - Decadal Predictability and Prediction

Climate Workshop Goals

(from Washington's presentation to BERAC)

- Review and identify the critical scientific challenges.
- Prioritize the challenges in terms of decadal or annual timelines.
- Identify the challenges where computing at the extreme scales is critical for climate change science success within the next two decades.
- Engage international scientific leaders in discussing opportunities to shape the nature of extreme scale scientific computing.
- Provide the high performance computing community with an opportunity to understand the potential future needs of the climate change research community.
- Look for breakthroughs.

Priority Research Directions (PRDs) were established for each of the Breakout sessions

Some PRDs are highlighted as follows:

PRDs for Model Development and Integrated Assessment

(from Washington's presentation to BERAC)

- How do the carbon, methane, and nitrogen cycles interact with climate change?
- How will local and regional water, ice, and clouds change with global warming?
- How will the distribution of weather events, particularly extreme events, that determine regional climate change with global warming?
- What are the future sea level and ocean circulation changes?

PRDs for Algorithms and Computational Environment

(from Washington's presentation to BERAC)

- Develop numerical algorithms to efficiently use upcoming petascale and exascale architectures
- Form international consortium for parallel input/output, metadata, analysis, and modeling tools for regional and decadal multimodel ensembles
- Develop multicore and deep memory languages to support parallel software infrastructure
- Train scientists in the use of high-performance computers.

PRDs for Decadal Predictability and Prediction
(from Washington's presentation to BERAC)

- Identify sources and mechanisms for potential decadal predictability
- Develop strategies for tapping into this predictability and ultimately realizing predictions that have societal benefit

PRDs for Data Visualization and Computing Productivity

(from Washington's presentation to BERAC)

- Develop new, robust techniques for dealing with the input/output, storage, processing, and wide-area transport demands of exascale data
- Integrate diverse and complex data
- Dedicate resources to the development of standards, conventions, and policies, and contribute to related committees

Office of Science
U.S. DEPARTMENT OF ENERGY

Diverse and complex data are integrated into visualizations to communicate model predictions

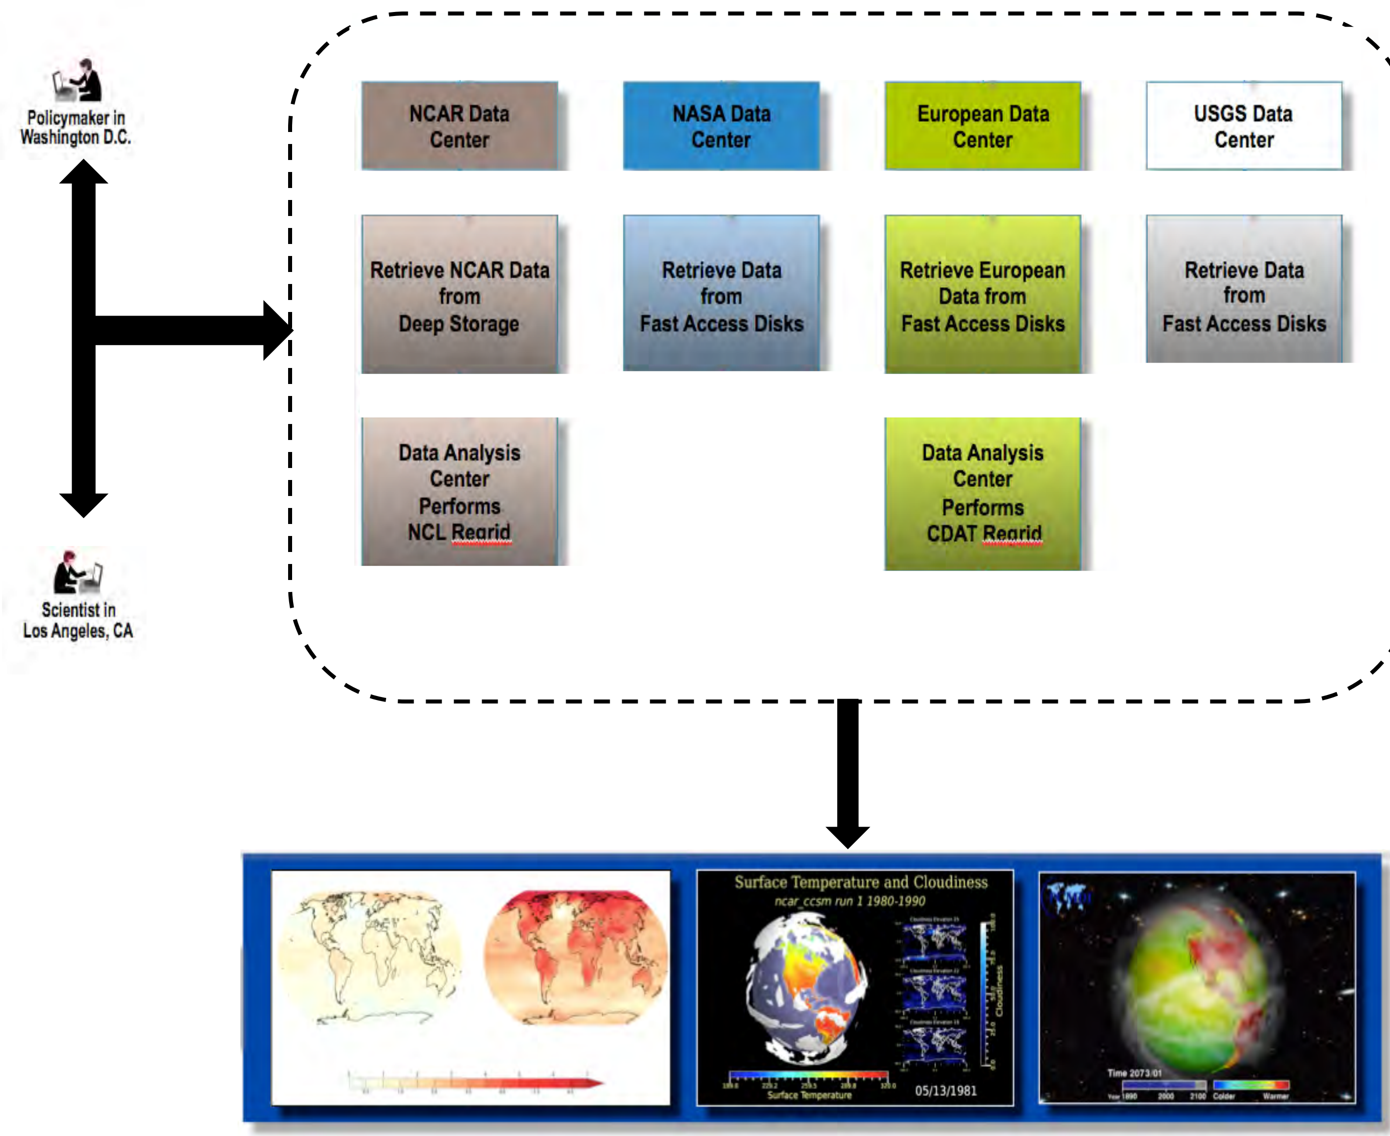

Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale

December 9-11, 2008, SLAC

- Workshop chair: Roger Blandford, co-chairs: Norman Christ, Young-Kee Kim
- Panel topics
  - Astrophysics data
  - Cosmology and astrophysics simulations
  - Experimental particle physics
  - Accelerator simulation
  - High energy theoretical physics

Selected PRDs identified by HEP workshop

- Cosmology and Astrophysics Simulation
  - Conduct cosmic structure formation probes of dark universe
  - Understand and calibrate supernovae as probes of dark energy
- Accelerator simulation
  - Design lepton collider linac module for cost and risk reduction
  - Predict beam loss and activation in Intensity Frontier accelerators and maximize Energy Frontier accelerators

Cosmic Structure Formation Probes of the Dark Universe

Scientific and computational challenges

Understand cosmic structure to enable the use the universe as a probe of fundamental physics

Perform cosmological hydrodynamical simulations with the dynamic range necessary to interpret future experiments

Summary of research direction

Develop precise predictions of structure formation from the Hubble Volume to the scale of the Solar System

Develop spatially and temporally adaptive codes, algorithms, and workflows for simulations and data on extreme-scale architectures.

Potential scientific impact

Determine the equation of state of dark energy and distinguish between dark energy and modifications of General Relativity

Measure the masses and interactions of dark matter

Measure the sum of the neutrino masses

Probe the fields responsible for primordial fluctuations

Potential impact on High Energy Physics

Revolutionize High Energy Physics by discovering and measuring physics beyond the standard model inaccessible to accelerators.

10 years

The Software Dimension
Consensus view of Astrophysics Simulation and Data Panels

- Identify and support development of low-level modules and libraries, isolating architectural complexity (e.g., MPI, FFT)
- Identify and support development of open-source community application codes, but not to the exclusion of other promising efforts
- Promote development of data models and language for interoperable data analysis (observation <=> simulation)

Forefront Questions in Nuclear Science and the Role of High Performance Computing

January 26-28, 2009, DC

- Workshop chair: Glenn Young,
  co-chairs David Dean, Martin Savage
- Panel topics
  - Nuclear structure and nuclear reactions
  - Nuclear astrophysics
  - Nuclear forces and cold QCD
  - Hot and dense QCD
  - Accelerator physics

Selected PRDs from NP workshop

- Physics of extreme neutron-rich nuclei and matter
  - which includes computing properties of nuclei that determine the r-process nucleo-synthesis path in stars and the nucleonic matter in neutron star cores and crusts.
- Microscopic description of nuclear fission
  - which involves a problem of both basic science interest and of great practical importance—computing the paths to fission and its products.
- Nuclei as neutrino physics laboratories
  - which involves computing properties of nuclei in double-beta decay experiments and neutrino-nucleus cross sections for supernova calculations and neutrino decays.
- Reactions that made us
  - which involves computing the triple-alpha process that produces ^{12}C, the nucleus at the core of organic chemistry and thus life, and $^{12}C(\alpha,\gamma)^{16}O$, the element that is key to both water and the reactions that power humankind and our present society.
- Each of these areas requires the solution of the nuclear quantum many-body problem and includes development and scaling of techniques that encompass two and three-nucleon forces, derived from QCD, with a strongly repulsive short-range character.

Nuclear Physics Workshop
Summary Observations
by Martin Savage

- Exa-scale computing resources are required to achieve the central mission of Nuclear Physics
  - theoretical AND experimental
  - science case is clear
- Nuclear Physics will be considerably unified (transformed) by Exa-scale computational resources
- Organizational infrastructure is required to enable full utilization of Exa-scale resources -- start organizing
- Collaborations within and external to Nuclear Physics are vital
- Prompt access to cutting edge machines is essential (- machine design)

Exa-scale Computational Resources

(slide courtesy Martin Savage)

- Meeting structured around present Nuclear Physics areas of effort

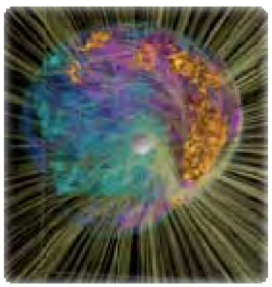

Nuclear Astrophysics

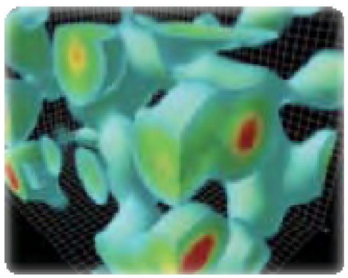

Cold QCD and Nuclear Forces

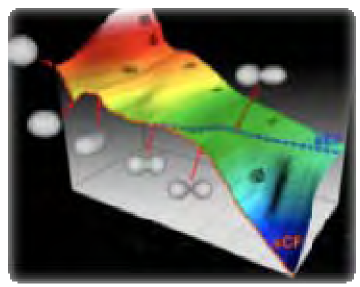

Nuclear Structure and Reactions

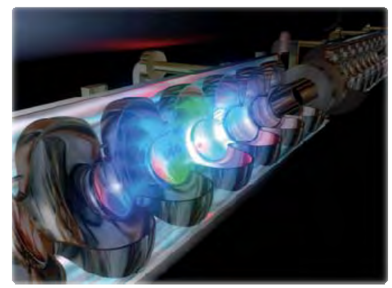

Accelerator Physics

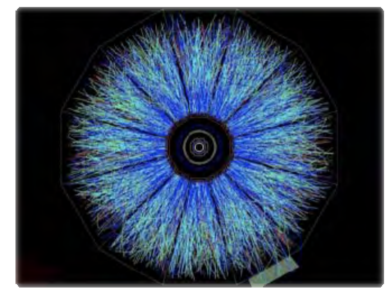

Hot and Dense QCD

- Exa-scale computing is REQUIRED to accomplish the Nuclear Physics mission in each area
- Staging to Exa-flops is crucial :
  - 1 Pflop-yr to 10 Pflop-yrs to 100 Pflop-yrs to 1 Exa-flop-yr (sustained)

Nuclear Physics Requires Exa-scale Computation

(slide courtesy Martin Savage)

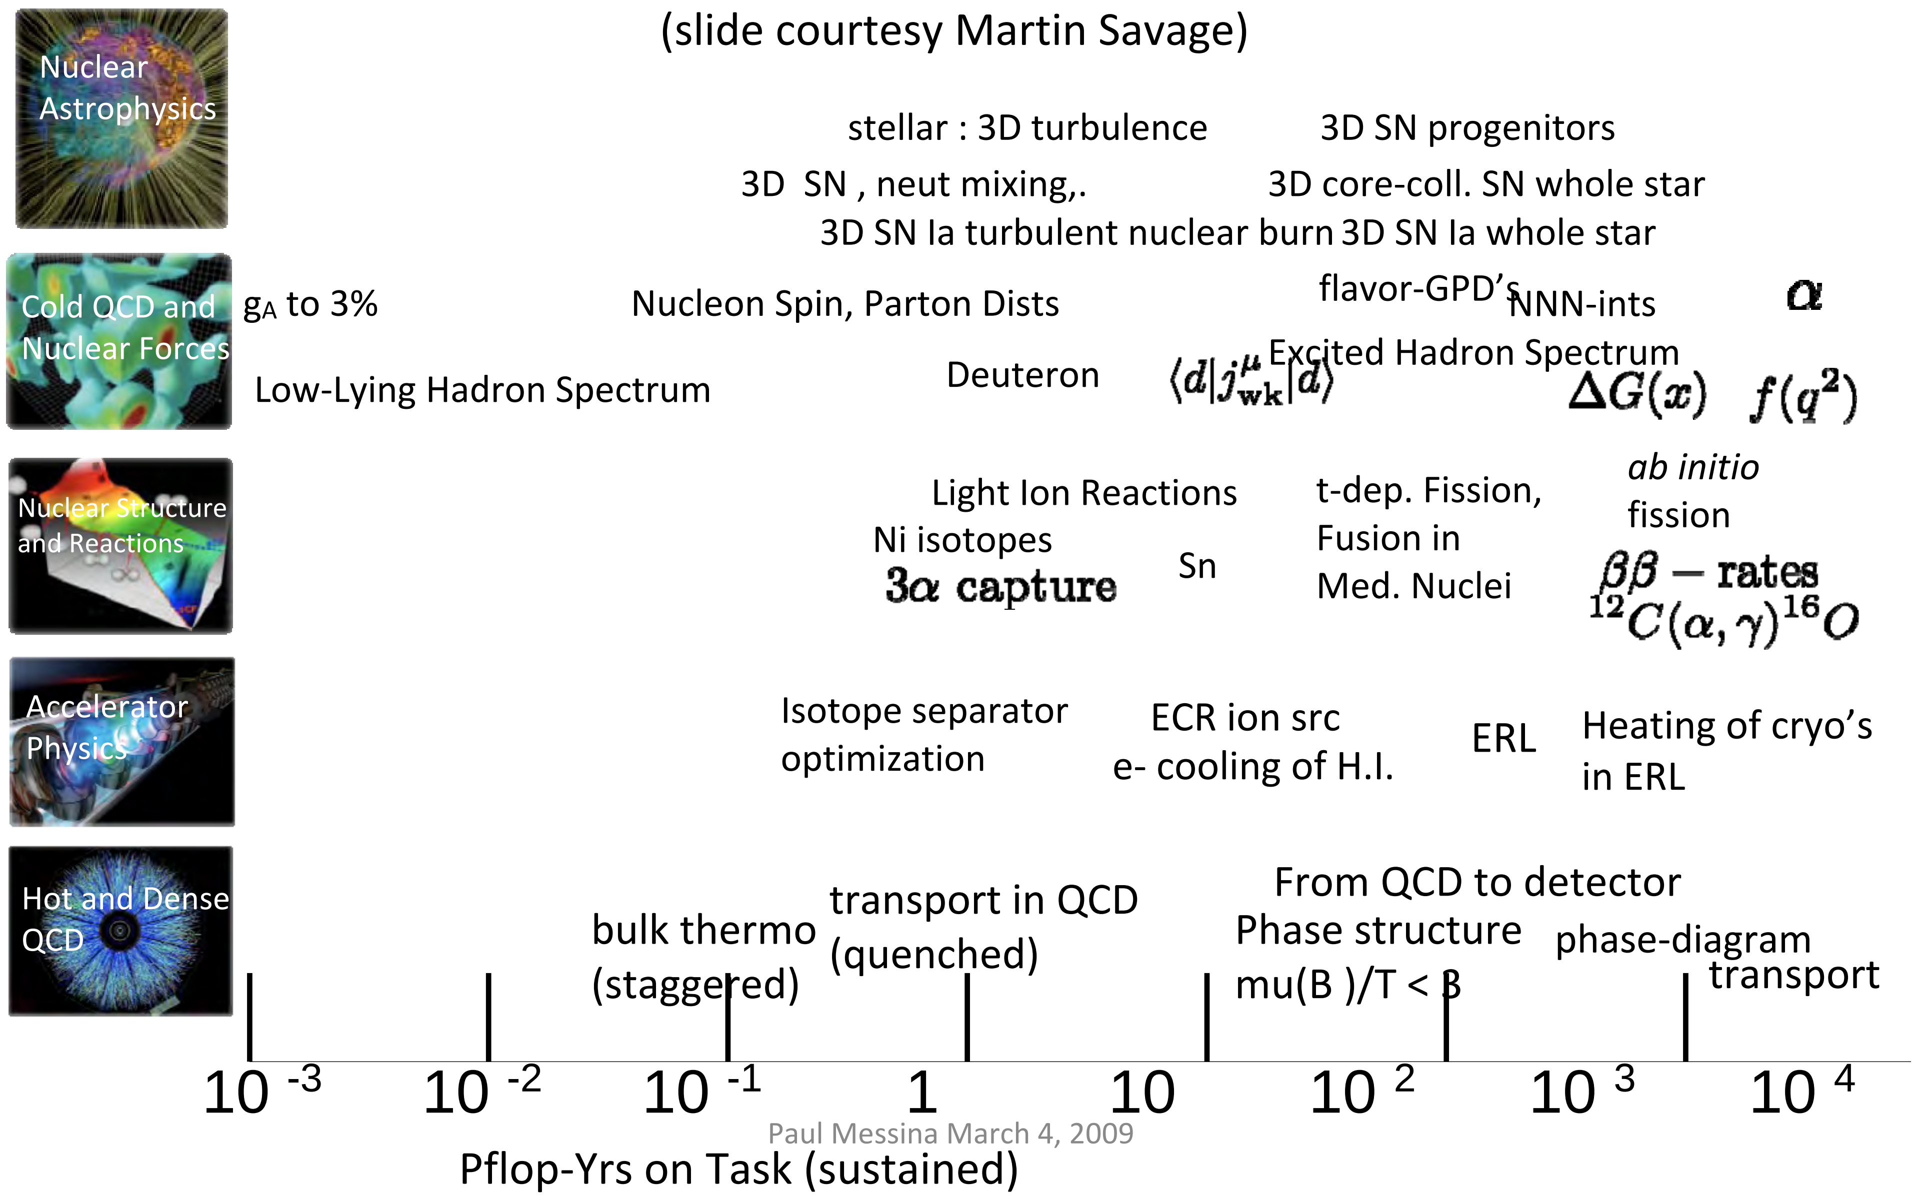

Exa-scale computing will unify Nuclear Physics

(slide courtesy Martin Savage)

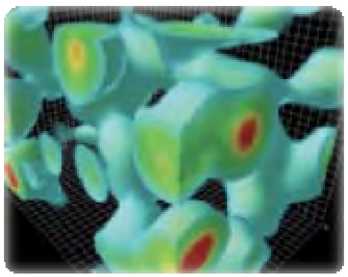

Cold QCD and
Nuclear Forces

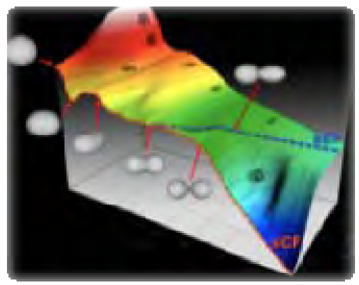

Nuclear Structure
and Reactions

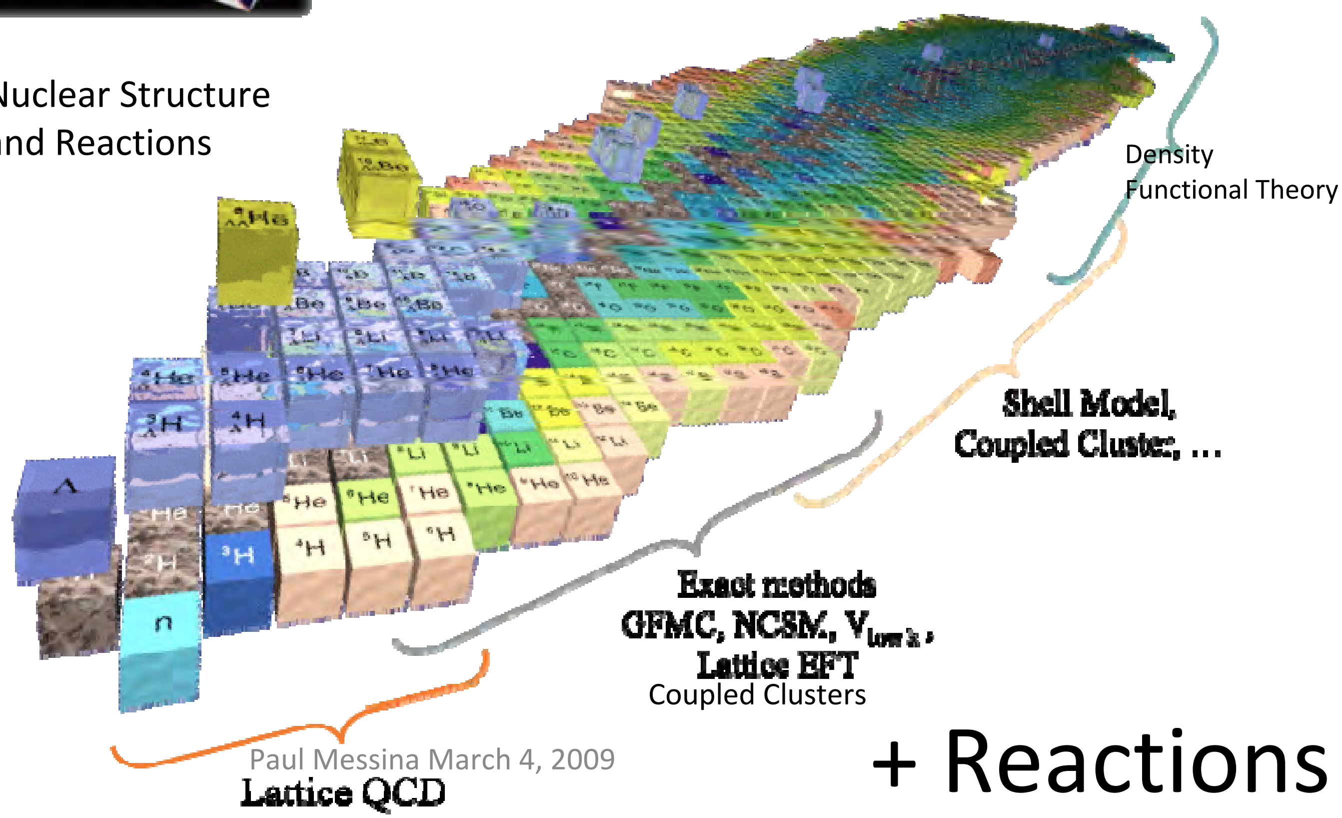

+ Reactions

Different Hardware Requirements

(slide courtesy Martin Savage)

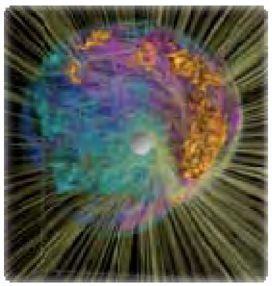

Nuclear Astrophysics

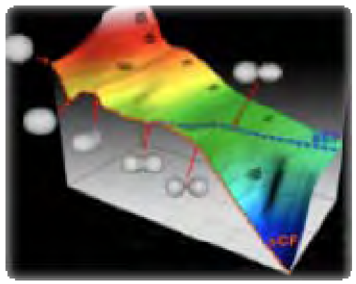

Nuclear Structure and Reactions

- Significant memory/core
- Fast memory access
- somewhat conflicts with present machine design ?
- low efficiency
- in current design : 4 GB/core preferred
- IO

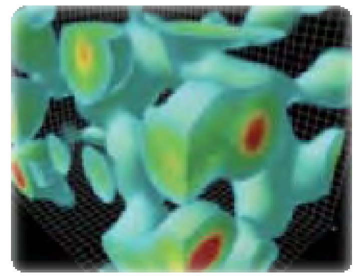

Cold QCD and Nuclear Forces

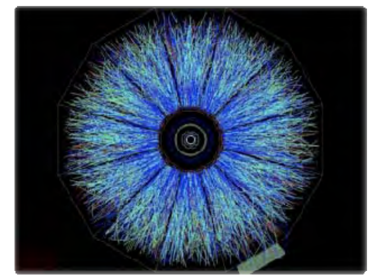

Hot and Dense QCD

- Modest memory per core
- Large number of cores
- Range of Partitions
- Fast network
- low latency
- IO

Scientific Grand Challenges in Fusion Energy Sciences and the Role of Computing at the Extreme Scale

March 18-20, 2009, DC

- Co-Chairs: Bill Tang and David Keyes
- Panel topics – Plasma and fusion energy sciences
  - Burning plasma/ITER
  - Advanced physics integration
  - Plasma-material interaction
  - Laser-plasma interactions
  - Magnetic reconnection physics

Scientific Grand Challenges in Fusion Energy Sciences and the Role of Computing at the Extreme Scale

March 18-20, 2009 DC

- Panel topics: Cross-cutting areas
  - Scalable algorithms
  - Data analysis, management, and visualization
  - Mathematical formulations
  - Programming models, frameworks, and tools (incl. languages, optimization, etc.)

Opportunities in Biology at the Extreme Scale of Computing

April 30-May 2, 2009, Chicago

- Co-Chairs: Mark Ellisman and Rick Stevens
- Panel topics
  - Tissues, Organs and Physiology Modeling
  - Pathways, Cells and Organelles
  - Proteins and Protein Complexes
  - Populations, Communities, Ecosystems and Evolutionary Dynamics
  - Genomics and metagenomics etc.
  - Imaging and Computer in the Loop
- Keynote speakers
  - David Kingsbury - The Gordon and Betty Moore Foundation - Palo Alto
  - Henry Markram - The Blue Brain Project - Lausanne/Geneva

Science Based Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale

May 11-12, 2009, DC

- Co-Chairs: Ernie Moniz and Bob Rosner
- Panel topics
  - Integrated Performance and Safety Simulations of Nuclear Energy System
  - Advanced Material Behavior Modeling
  - Verification, Validation and Uncertainty Quantification for Nuclear Energy Simulations
  - Nuclear Energy System Integration

Discovery in Basic Energy Sciences: The Role of Computing at the Extreme Scale

August 12-14, 2009, DC

- Co-Chairs: Giulia Galli and Thom Dunning
- Preliminary panel topics
  - Combustion
  - Many Body Effects, Correlation and Excited States: QMC, DMFT, Conventional QC
  - Catalysis & Electrochemistry
  - Time Dependent Processes
  - Nanoscale Materials and Technology

Scientific challenges for the NNSA and Office of Science missions

September or October 2009, DC

- Co-Chairs: Alan Bishop and Paul Messina
- Panel topics:
  - HYDRODYNAMICS
    - rad-hydro, complex fluids, turbulent flow/shocks, plasmas ....
  - MATERIALS SCIENCE
    - EOS, aging, fracture, friction, plasticity, ejecta, multi-phase/composites, properties under extreme conditions ....
  - CHEMISTRY/BIOLOGY
    - combustion, conventional explosives, polymers/foams, catalysis, surface chemistry, quantitative biology, quantum simulations of atomic/molecular structure/response …
  - NUCLEAR/PARTICLE PHYSICS
    - nuclear structure/cross-sections, fission, astrophysics/cosmology, nuclear energy …
  - INFRASTRUCTURE/SYSTEM/NETWORK SIMULATIONS
    - energy-climate, power grid, disease spread, cognition .....
  - VERIFICATION, VALIDATIONS, AND UNCERTAINTY QUANTIFICATION/ERROR ANALYSIS
    - Separate panel as well as included in all other panels

Related Workshops

- We will organize workshops aimed at defining the properties of the computing environments that will enable researchers to tackle the science challenges, guided by the findings of the eight "Science Challenge" workshops
- The workshops that the International Exascale Software Project is organizing will identify the key software environment and tools that are necessary for productive use of present and future leadership computer systems and viable approaches to developing them

Credit where credit is due

- Michael Strayer had the vision to conceive these workshops and the other DOE-SC Associate Directors have given their support
- Walt Polansky and Lali Chatterjee have worked with the other DOE program offices to create a vision for the workshops and provided guidance on their scope and organization
- Moe Khaleel and his PNNL team provide invaluable support and contributions at all levels
- The community members who have organized and participated in the workshops are identifying exciting science challenges that could be tackled with future computer environments of much greater power than today's