U.S. Department of Energy

Office of Science

Modeling and Simulation at the Exascale for Energy and the Environment (E3)

Horst Simon (LBNL), Rick Stevens (ANL),

Thomas Zacharia (ORNL)

ASCAC Meeting

Washington D.C., August 15, 2007

U.S. Department of Energy

Office of Science

Modeling and Simulation at the Exascale for Energy and the Environment

U.S. Department of Energy
Office of Science

Simulation and Modeling at the Exascale for Energy, Ecological Sustainability and Global Security An Initiative

The objective of this ten-year vision, which is in line with the Department of Energy's Strategic Goals for Scientific Discovery and Innovation, is to focus the computational science experiences gained over the past ten years on the opportunities introduced with exascale computing to revolutionize our approaches to energy, environmental sustainability and security global challenges.

Executive Summary

The past two decades of national investments in computer science and high-performance computing have placed the DOE at the forefront of many areas of science and engineering. This initiative capitalizes on the significant gains in computational science and boldly positions the DOE to attack global challenges through modeling and simulation. The planned petascale computer systems and the potential for exascale systems shortly provide an unprecedented opportunity for science; one that will make it possible to use computation not only as an critical tool along with theory and experiment in understanding the behavior of the fundamental components of nature but also for fundamental discovery and exploration of the behavior of complex systems with billions of components including those involving humans.

Through modeling and simulation, the DOE is well-positioned to build on its demonstrated and widely-recognized leadership in understanding the fundamental components of nature to be a world-leader in understanding how to assemble these components to address the scientific, technical and societal issues associated with energy, ecology and security on a global scale.

For these types of problems, the time-honored, or subsystems, approach in which the forces and the physical environments of a phenomenon are analyzed, is approaching a state of diminishing returns. The approach for the future must be systems based and simulation programs are developed in the context of encoding all known relevant physical laws with engineering practices, production, utilization, distribution and environmental factors.

This new approach will

- **Integrate, not reduce.** The full suite of physical, chemical, biological, chemical and engineering processes in the context of existing infrastructures and human behavior will be dynamically and realistically linked, rather than focusing on more detailed understanding of smaller and smaller components.
- **Leverage the interdisciplinary approach to computational sciences.** Current algorithms, approaches and levels of understanding may not be adequate. A key challenge in development of these models will be the creation of a framework and semantics for model

The objective of this ten-year vision, which is in line with the Department of Energy's Strategic Goals for Scientific Discovery and Innovation, is to focus the computational science experiences gained over the past ten years on the opportunities introduced with exascale computing to revolutionize our approaches to energy, environmental sustainability and security global challenges.

Based on this initial white paper, ANL, LBNL, and ORNL organized the community input process in the form of three town hall meetings.

The Opportunity

- Attack global challenges through modeling and simulation
- Planned petascale and the potential exascale systems provide an unprecedented opportunity
- Beyond computation as an critical tool along with theory and experiment
- Understanding the behavior of the fundamental components of nature
- Fundamental discovery and exploration of complex systems with billions of components including those involving humans

Town Hall Meetings

April 17-18 at LBNL

http://hpcrd.lbl.gov/E3SGS/main.html

May 17-18 at ORNL

http://computing.ornl.gov/workshops/town_hall/index.shtml

May 31-June 1 at ANL

https://www.cls.anl.gov/events/workshops/townhall07/index.php

About 450 (unique) participants

Goals for Town Hall Meetings(1)

- **To gather community input for possible future DOE research initiatives in the areas of high-performance computing, computer science, computational science and advanced mathematics and the application of these to global challenge problems**

- **To examine the prospects for dramatically broadening the reach of HPC in to new disciplines, including areas such as predictive modeling in biology and ecology, integrative modeling in earth and economics sciences, bottoms up design for energy and advanced technologies**

Goals (2)

- To identify emerging domains of computation and computational science that could have dramatic impact on economic developments such as agent-based simulation, self-assembly and self-organization
- To outline the challenges and opportunities for Exascale capable systems, ultra low power architectures and ubiquitous multi-core technologies (inc software, etc.)
- To identify new opportunities for end-to-end investment in new computational science problem areas (including validation and verification)

Break Out Groups (applications)

B1. **Improve our understanding of complex biogeochemical (C, N, P, etc.) cycles that underpin global ecosystems functions and control the sustainability of life on Earth.**

B2. **Develop and optimize new pathways for renewable energy production and development of long-term secure nuclear energy sources, through computational nanoscience and physics-based engineering models.**

B3. **Enhance our understanding of the roles and functions carried out by microbial life on Earth, and adapt these capabilities for human use, through bioinformatics and computational biology.**

B6. **Develop integrated modeling environments that couple the wealth of observational data and complex models to economic, energy, and resource models that incorporate the human dynamic into large-scale global change analysis.**

B9. **Develop a "cosmic simulator" capability that integrates increasingly complex astrophysical measurements with simulations of the growth and evolution of structure in the universe, linking the known laws of microphysics to the macro world. Develop large-scale, special-purpose computing devices and innovative algorithm development to achieve this goal.**

B10. **Manufacturing**

Break Out Groups (technology)

B4. **Develop tools and methods to protect the distributed information technology infrastructure: ensuring network security, preventing disruption of our communications infrastructure, and defending distributed systems against attacks.**

B5. **Drive innovation at the frontiers of computer architecture and information technology, preparing the way for ubiquitous adoption of parallel computing, power-efficient systems, and the software and architectures needed for a decade of increased capabilities. Accelerate the development of special-purpose devices that have the potential to change the simulation paradigm for certain science disciplines.**

B7. **Advance mathematical and algorithmic foundations to support scientific computing in emerging disciplines such as molecular self- assembly, systems biology, behavior of complex systems, agent-based modeling and evolutionary and adaptive computing.**

B8. **Integrate large, complex, and possibly distributed software systems with components derived from multiple applications domains and with distributed data gathering and analysis tools.**

The Charge Questions

1. What (in broad brush) is feasible or plausible to accomplish in 5-10 years?
2. What are the major challenges in the area?
3. What is today's state-of-the art in the area?
4. How would we accelerate development?
5. What are expected outcomes and impact of acceleration or increased investment (i.e., what problems would we aim to solve or events we would cause to occur)?
6. What scale of investment would be needed to accomplish the outcome?
7. What are the major risks?
8. What and who are missing?

U.S. Department of Energy

Office of Science

Climate/Biogeochemical Modeling - Findings

How can we improve our understanding of complex biogeochemical cycles that underpin global ecosystems functions and control the sustainability of life on Earth?

- **Improvements in representation of biogeochemical cycles in ESMs can be achieved by using a combination of data assimilation and development of mechanistic and process-based models.**
- **Higher spatial resolution is needed to address the fine-scale heterogeneity inherent in biogeochemical process.**
- **New, innovative approaches are needed both in fundamental applied mathematics and in computational science to quantify the uncertainty inherent in a large systems-level model such as the ESM.**

Climate/Biogeochemical Modeling - Challenges

To develop process-scale biological and ecological process modeling within the Earth system, to develop new methodologies and software tools that integrate observations into these branches of Earth system science and advance existing ESMs, and to quantify uncertainties at regional to local scales.

- **High-resolution ESMs with massive assimilation of satellite and other data.**
- **Detailed modeling of controlled and modified ecosystems to fit the environmental envelope in which future climate changes will occur.**
- **Development of process-scale mechanistic models for biogeochemical, hydroecological, cloud microphysical, and aerosol processes.**
- **Rational design and analysis of computer experiments to navigate very large parameter space with very large outputs.**

Energy – Findings

Providing new models and computational tools with the functionality needed to discover and develop complex processes inherent in a new energy economy

Three pathways to a low-carbon economy: computational nanoscience and materials science for renewable energy; simulation modeling for a fusion pathway: and simulation and modeling for advance nuclear energy systems.

- **Computational Nanoscience and Material Science for Renewable Energy**
  - Needs materials optimized for hydrogen storage
  - Needs reliable and efficient catalysis for water dissociation in hydrogen production
  - Needs cost-effective, environmentally benign, and stable material for efficient solar cells
- **Advanced Nuclear Energy Systems**
  - Spent fuel reprocessing is very complicated and requires a large number of different materials - multiple pathways must be considered
  - Waste streams must be treated
  - Improved coupling between computations and experiments must occur
- **Fusion Energy - the promise of ITER**
  - Designed to produce 500 million Watts of heat from fusion reactions for over 400 seconds with gain exceeding 10 – thereby demonstrating the scientific and technical feasibility of magnetic fusion energy
  - Fusion fuel will be sustained at high temperature by the fusion reactions themselves
  - Data from experiments worldwide, supported by advanced computation, indicate that ITER is likely to achieve its design performance

Energy - Challenges

- **Computational Nanoscience and Material Science for Renewable Energy**
  - The possible exploratory parameter spaces are huge
  - Potential for device improvement - tremendous challenge to find the best material and design
  - Optimization with exascale computation using a direct numerical material by design search, or by understanding some fundamental processes in nanosystems
- **Advanced Nuclear Energy Systems**
  - The main challenge is to significantly reduce the radioactive waste; want a nuclear fuel cycle with a factor of a hundred less long-lived waste
  - Reduce fuel development and qualification time
  - Assess life cycle performance
  - Address safety concerns
  - Predict fuel rod behavior in design basis accident
  - Predict transuranic fuel behavior
  - Eliminate unrealistic assumptions that drive to more complex designs and thus higher installation cost
  - Achieve higher power efficiencies
  - Reduce the learning curves to get efficiencies
  - Reduce the required number of repositories
- **Fusion Energy - the promise of ITER**
  - Science and technology is needed to achieve the continuous power with increased gain in a device of similar size and field
  - Strong R & D programs are needed to harvest the scientific knowledge from ITER and leverage its results
  - Advanced computations in tandem with experiment and theory are essential
  - Accelerated development of computational tools and techniques that aid the acquisition of the scientific understanding needed to develop predictive models which can prove superior to extrapolations of experimental results

Microbial Life - Findings

- **Microbes have been on earth for at least 3.5 Billion years, they are responsible for the O_2 in the atmosphere, the N that enables plant growth and sequestering much of the CO_2 in the oceans**
- **Microbial processes lie at the base of all ecosystem functions on earth and provide about 50% of the primary productivity of the biosphere and 98% of the primary productivity of oceans**
- **Microbial diversity is huge: 10^7 species and 10^{30} cells, 10^{36} microbial proteins on the planet and 10,000 times the surface area of the Earth**
- **There are hundreds of useful molecular machines (proteins) that have been "mined" from microbes and probably thousands of useful machines yet to be discovered.**
- **Computational techniques are now enabling us to begin to understand this diversity through bioinformatics**
  - Reconstructing their genomes and identifying novel proteins
  - Modeling their metabolisms
  - Modeling their complex multispecies communities

Microbial Life - Challenges

- **Model driven HT experimental data generation**
- **Improving Model development**
  - Genome scale metabolic networks, regulatory networks, signaling and developmental pathways
  - Microbial ecosystems and complex biogeochemical interactions
- **Bioinformatics techniques to address the integration of genomics, proteomics, metagenomics and structural data to screen for novel protein function discovery**
- **Molecular modeling techniques that can address the multiscale challenges**

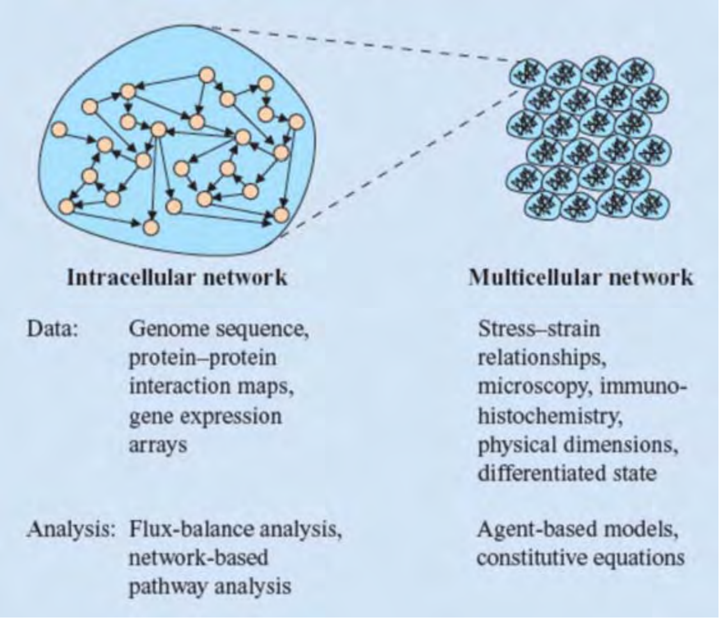

An Integrated View of Modeling, Simulation, Experiment, and (Bio)informatics

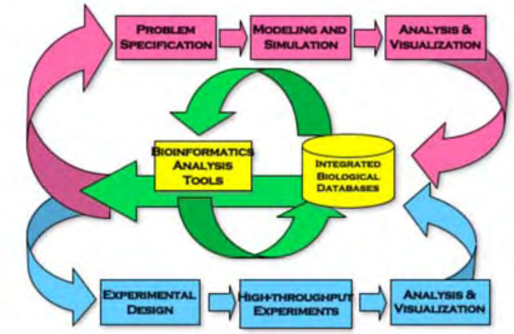

Socio-Economic Modeling: Findings

- **We must model human responses to climate change if we are to understand their likely effectiveness and impacts, and thus help to sustain a prosperous and secure society.**
- **Exascale computers have the potential to transform understanding of socio-economic-environmental interactions through more detailed treatments of the various components and their interactions, and of issues of uncertainty and risk.**
- **Global socio-economic modeling is ripe for Exascale computing.**
- **Opportunities include the creation of coupled socio-economic-environmental models and greatly improved statistical analysis of data.**
- **Substantial effort must also be devoted to data quality and parameter estimation issues.**

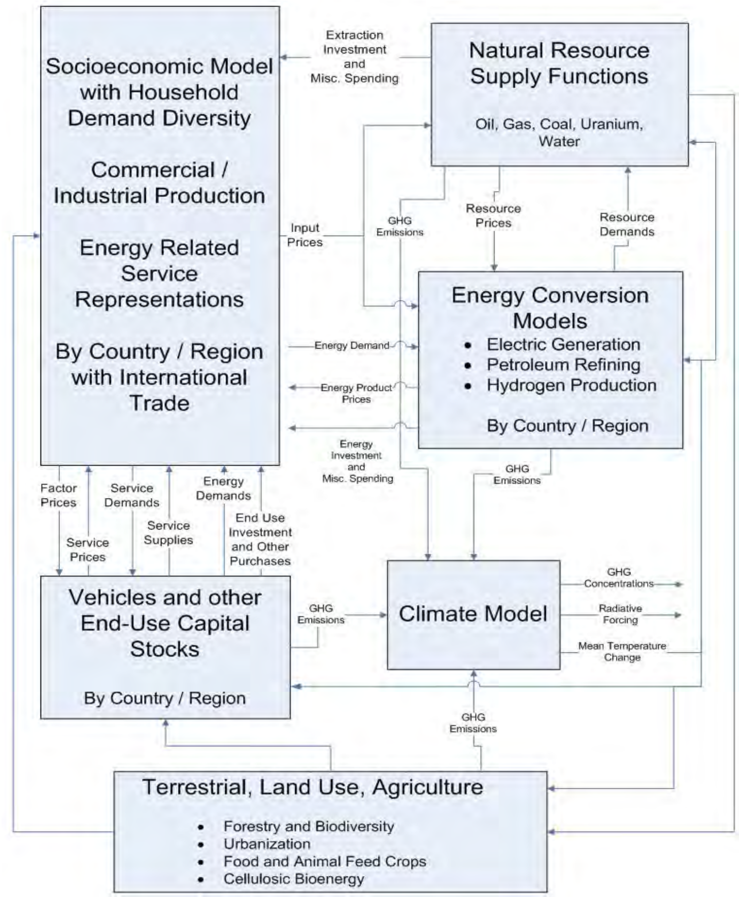

Socio-Economic Modeling: Challenges

- Construction of a comprehensive suite of models of unprecedented geospatial and temporal detail with comprehensive error analysis on the representation.
- Leverage state-of-the-art climate modeling activities (e.g., SCIDAC) to include economic prediction models under alternative climate regimes.
- Basic research into spatial statistics, modeling of social processes, relevant micro-activity and biosphere coupling issues, and relevant mathematical challenges, such as multiscale modeling.
- Assembly and quality control of extensive data collections
- Comprehensive and detailed validation of both individual models and large model systems.
- Development of novel, robust numerical techniques and high-performance computing approaches to deal with the expected orders-of-magnitude increase in model complexity.
- A wide range of application studies aimed at both validation and application.
- Education programs aimed at training the next generation of computational economists and other social scientists, including not only formal training programs but also web-based modeling and simulation tools that allow widespread access to the new models and their results.

U.S. Department of Energy

Office of Science

Astrophysics: Findings and Opportunities at Exascale

Area	Science
Large Scale Structure Formation	Simulations of the large-scale distribution of galaxies and galaxy clusters over a large fraction of the observable universe with one percent precision, required of the observational program proposed by the DOE/NASA/NSF-sponsored Dark Energy Task Force.
Galaxy Formation	Simulations of galaxy formation with sufficient resolution to predict the observed properties of individual galaxies in a volume containing a sufficient fraction of the observable universe to compare with large-scale surveys. Simulations of the formation of the Milky Way Galaxy with sufficient precision to compare with data from JWST and LSST.
Stellar Evolution	Simulations of the entire stellar envelope in AGB stars, responsible for the supply of half of the heavy elements (elements above iron) in Nature.
Supernovae	Definitive 3D multiphysics simulations of core collapse supernovae, the dominant source of elements between oxygen and iron and the half of the heavy elements not produced in AGB stars.
Compact Objects	Definitive simulations of binaries involving two neutron stars, and one black hole and one neutron star, which are among the leading candidates for the production of gravitational waves in our galaxy.

Astrophysics: Challenges

- Increased use of adaptive mesh refinement will require increased need for dynamic load balancing
- Long run times (even) on exascale computers will require increased need for fault tolerance
- Large volume of experimental data will require new data analysis and visualization techniques

Mathematics and Algorithms Findings

The current state-of-the-art is characterized in the following areas:

- Solvers
- Uncertainty Quantification
- Adaptive Mesh Refinement (AMR)
- Analysis in High Dimensional Spaces
- Data Analysis
- High Precision Arithmetic

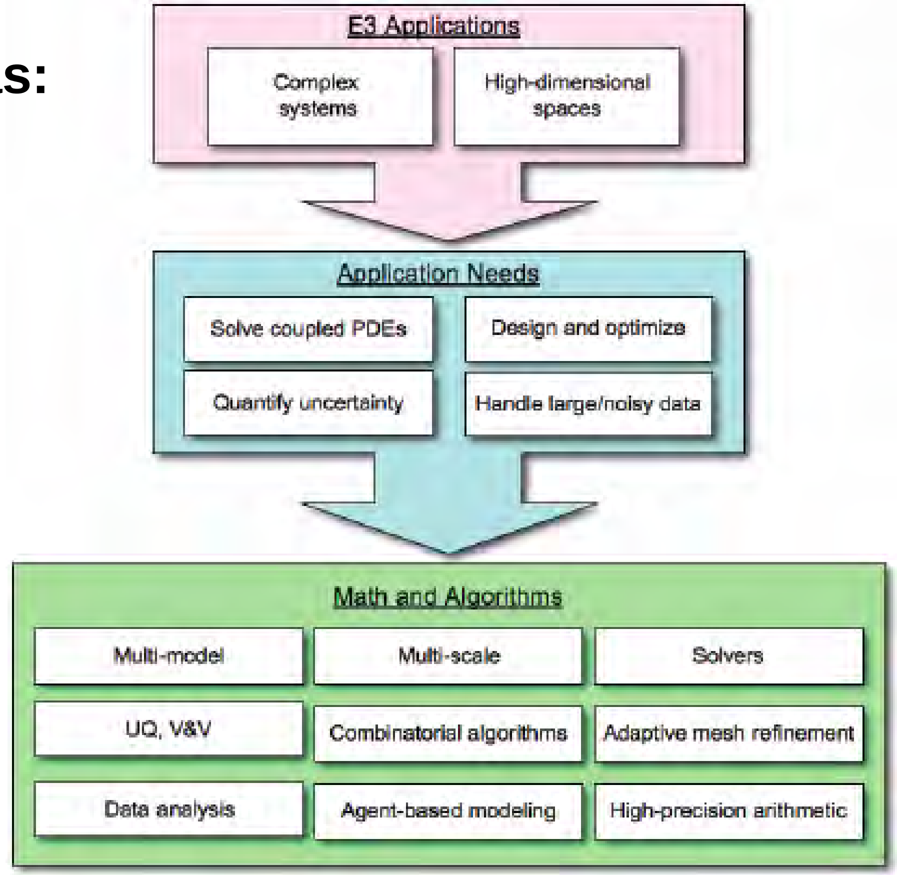

Mathematics and Algorithms Challenges

Advances in four interlinked areas will be required:

- **Coupled Models:** multiple models (PDE or data-based) for different phenomena or at different scales; scalable implementations for coupled systems; extend existing models and codes in terms of scalability; new mathematical approaches (general and domain-specific) to model coupling; new implicit approaches for dealing with long time-scale coupled simulations
- **Uncertainty** systematic approach for quantifying, estimating and controlling the uncertainty caused, for example, by reduced models, uncertain parameters, or discretization error; tools that automatically construct representations of uncertainty, handle the uncertainty propagation and coupling effects, and provide sharp estimates of the uncertainty of key merit criteria
- **Optimization** scalable algorithms for continuous nonlinear optimization; parallel branch-and-cut methods for linear and nonlinear optimization problems with discrete variables;more sophisticated parallel methods for solving stochastic optimization problems
- **Large datasets** new data representations, data handling algorithms, efficient implementations of data analysis algorithms on high-performance computing platforms, and representations of analysis results

Software - Findings

"Effective use of exascale systems will require fundamental changes in how we develop and validate simulation codes for these systems and how we manage and extract knowledge from the massive amount of data produced."

- **Exascale computer architectures necessitate radical changes to the software used to operate them and the science applications. The change is as disruptive as the shift from vector to distributed memory supercomputers 15 years ago.**

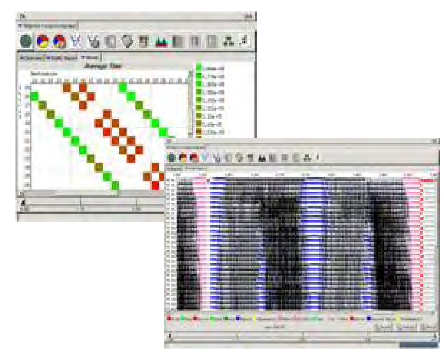

- **Message passing coupled with sequential programming languages will be inadequate for architectures based on many-core chips.**
- **Present code development, correctness, and performance analysis tools can't scale up to millions of threads.**
- **Checkpointing will be inadequate for fault tolerance at the exascale.**

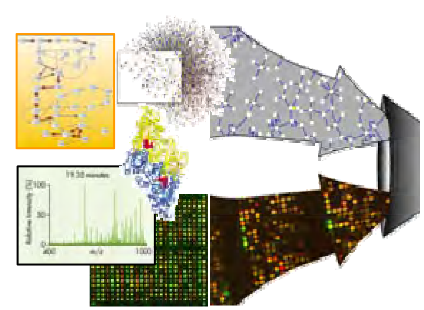

- **Fundamental changes are necessary to manage and extract knowledge from the tsunami of data created by exascale applications.**

Software - Challenges

Improve scientists' and administrators' productivity

- Creation of development and formal verification tools integrated with exascale programming models

Improve the robustness and reliability of the system and the applications.

- New fault tolerance paradigms will need to be developed and integrated into both existing and new applications

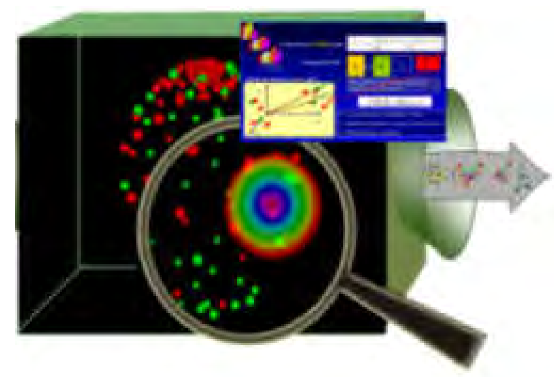

Integrate knowledge discovery into the entire software life-cycle

- Application development tools, runtime steering, post-analysis, and visualization

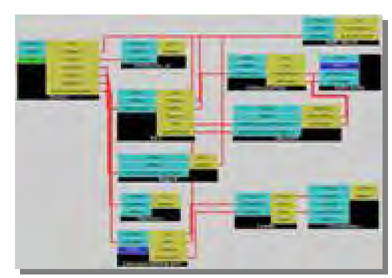

Develop new approaches to handling the entire data life-cycle of exascale simulations

- Seamlessly integration into the scientist's workflow
- Automatically capture provenance
- Develop effective formats for storing scientific data

Cyberinfrastructure – Findings

Even for the most basic workflow, such as collecting data, performing a high-performance computation, and analyzing the resulting data, it is rare that all of the necessary people, computing, storage, and analysis systems are within the same location

- **Workflow management;**
  - Workflow systems today are evolving at a promising rate but show no signs yet of coalescing around a compact set of solutions.
- **Collaboration frameworks and techniques;**
  - Asynchronous collaboration infrastructure today includes of the use of wikis, blogs, and other emerging social networking tools, whereas synchronous collaboration infrastructure include context and location-aware persistent visualization and collaboration environments
- **Data management and movement of exascale datasets and data collections;**
  - Many researchers today who are creating terascale and petascale datasets find that they spend a significant portion of their time managing data rather than scientific investigation.
- **Authorization and authentication for flexible interdisciplinary computational science teams ("virtual organizations");**
  - Currently, cyberinfrastructure management tools for the largest HPC systems are mostly collections of scripts and tools not well suited to managing these systems. and
- **Management tools, techniques, and methodologies to understand performance of the infrastructure and to protect it from attack, disruption, and data loss.**
  - Most organizations protect their information technology resources through a defense-in-depth approach that covers network, host and application security technologies, provides cyber security awareness, skills development and training to staff and users and details cyber security policy and standards.

Cyberinfrastructure - Challenges

- **Representing information and reducing information** overload in a natural, usable, collaborative manner so that it is understandable and accessible to the researchers involved.
- **Scalable, flexible, and federated approach to authentication and authorization** and to the creation and management of the virtual organizations that manage collaboration resources
- **Higher performance data management and movement tools** and techniques must be developed for data centers and archives, portals, and intersite and intrasite file transfers.
- **Tools and techniques for the configuration, verification, troubleshooting, and management** of complex systems
- Exascale resources will not have commercial security products available that scale to line rates or capacities for months or even years after the exascale resources are deployed
- A key challenge to cyber security methodologies and tools of the future will be the creation of a **framework and semantics for integrating information** in the individual cyber security component systems for situational awareness, anomaly detection, and intrusion response.
- Tools for data transfer must be developed that use **dedicated channels** to separate data from control communication and facilitate the application of graded levels of control for exascale. Control sessions represent the primary threat.

Many challenges are scale independent, but some are more acute at the Exascale

Advanced Architectures Findings

- Exascale systems are likely feasible by 2017±2
- 10-100 Million processing elements (cores or mini-cores) with chips perhaps as dense as 1,000 cores per socket, clock rates will grow more slowly
- 3D packaging likely
- Large-scale optics based interconnects
- 10-100 PB of aggregate memory
- > 10,000's of I/O channels to 10-100 Exabytes of secondary storage, disk bandwidth to storage ratios not optimal for HPC use
- Hardware and software based fault management
- Simulation and multiple point designs will be required to advance our understanding of the design space
- Achievable performance per watt will likely be the primary measure of progress

Advanced Architectures Challenges

- **Performance per watt — stretch goal 100 GF/watt of sustained performance ⇒ 10 MW Exascale system**
  - Leakage current dominates power consumption
  - Active power switching will help manage standby power
- **Large-scale integration — need to package 10M-100M cores, memory and interconnect**
  - 3D packaging likely, goal of small part classes/counts
- **Heterogeneous or Homogenous cores?**
  - Mini cores or leverage from mass market systems
- **Reliability — needs to increase by 10^3 in faults per PF to achieve MTBF of 1 week**
  - Integrated HW/SW management of faults
- **Integrated programming models (PGAS?)**
  - Provide a usable programming model for hosting existing and future codes

U.S. Department of Energy
Office of Science

Summary Report