Scientific Grand Challenges for National Security:

THE ROLE OF COMPUTING AT THE EXTREME SCALE

October 6-8, 2009 • Washington D.C.

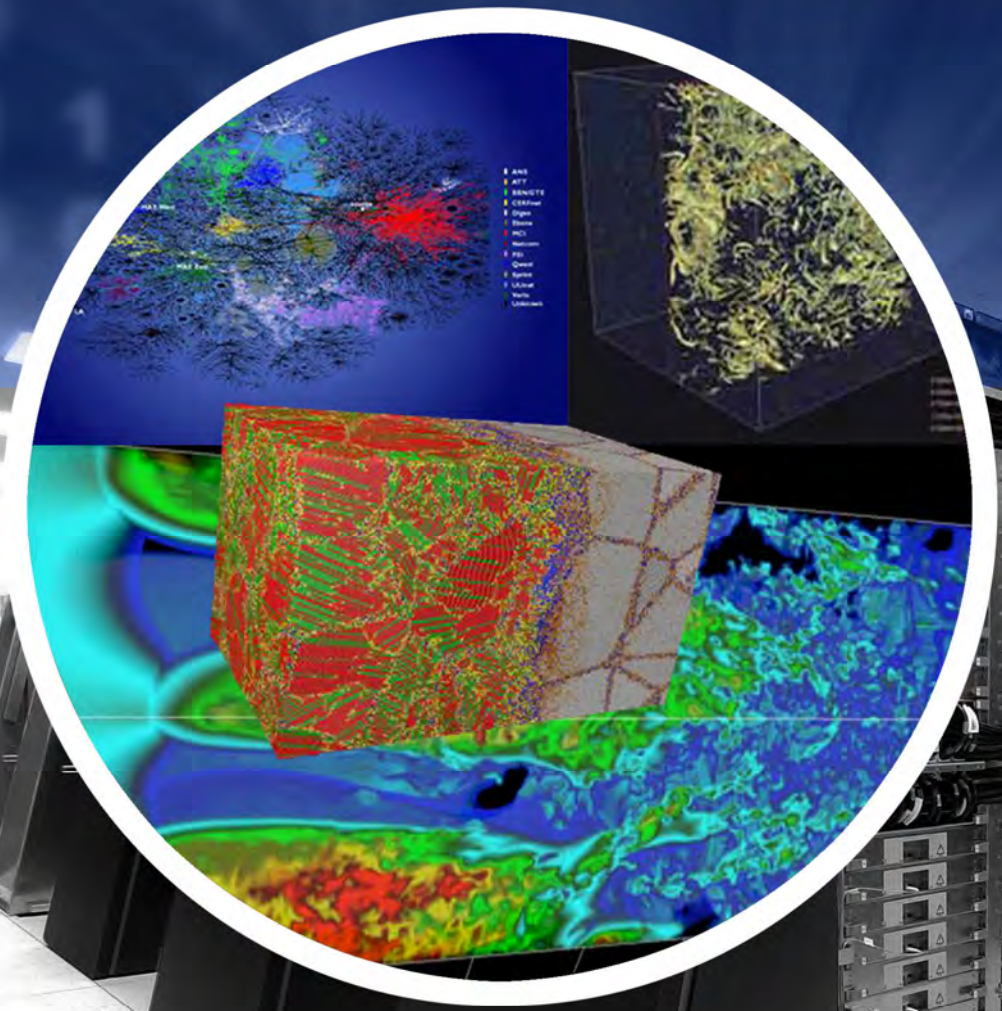

U.S. DEPARTMENT OF ENERGY

Sponsored by the Office of Advanced Scientific Computing Research and the National Nuclear Security Administration

DISCLAIMER

On the cover: The IBM Blue Gene/P supercomputer, dubbed the Intrepid, at the U.S. Department of Energy's Argonne National Laboratory. When the computer was installed in 2008, it was the fastest supercomputer in the world available to open science and the third fastest among all supercomputers. Future reports in the Scientific Grand Challenges workshop series will feature different Office of Science computers on their covers.

SCIENTIFIC GRAND CHALLENGES IN NATIONAL SECURITY: THE ROLE OF COMPUTING AT THE EXTREME SCALE

Report from the Workshop Held October 6-8, 2009

Sponsored by the U.S. Department of Energy, Office of National Nuclear Security Administration, and the Office of Advanced Scientific Computing Research

Chair, **Alan Bishop**
Los Alamos National Laboratory

Chair, **Paul Messina**
Argonne National Laboratory

Co-Lead, Multiphysics Modeling, **Robert Rosner**
University of Chicago

Co-Lead, Multiphysics Modeling, **Michael Zika**
Lawrence Livermore National Laboratory

Co-Lead, Nuclear Physics, **Joseph Carlson**
Los Alamos National Laboratory

Co-Lead, Nuclear Physics, **Edward P. Hartouni**
Lawrence Livermore National Laboratory

Co-Lead, Materials Science, **John Sarrao**
Los Alamos National Laboratory

Co-Lead, Materials Science, **Sidney Yip**
Massachusetts Institute of Technology

Co-Lead, Chemistry, **Robert Harrison**
Oak Ridge National Laboratory

Co-Lead, Chemistry, **Steve Plimpton**
Sandia National Laboratories

Co-Lead, The Science of Nonproliferation, **Mark D. Rintoul**
Sandia National Laboratories

Co-Lead, The Science of Nonproliferation, **Alex Pothen**
Purdue University

Co-Lead, Uncertainty Quantification and Error Analysis, **Dave Higdon**
Los Alamos National Laboratory

Co-Lead, Uncertainty Quantification and Error Analysis, **Richard Klein**
University of California, Berkeley, and Lawrence Livermore National Laboratory

---

Representative, Office of National Nuclear Security Administration, **Njema Frazier**
Representative, Office of National Nuclear Security Administration, **Robert Webster**
Representative, Office of Advanced Scientific Computing Research, **Lucy Nowell**
Representative, Office of Advanced Scientific Computing Research, **Lali Chatterjee**

EXECUTIVE SUMMARY

The last few decades have truly marked the beginning of a "century of complexity" in almost every scientific field that is relevant to the U.S. Department of Energy's (DOE) national security mission—from biology to materials to cosmology. Tackling the study of extremely complex phenomena is enabled by the accelerated pace of advancements in available measured and observed data, theoretical and modeling techniques for complex nonlinear phenomena, and equally impressive advances in simulation and computational resources. The observed evidence of functional, multiscale complexity in many disciplines and applications raises fundamental questions about the origins of this complexity, how to characterize it, and how to harness its consequences. The opportunity to move beyond better observation, measurement, and simulation capabilities to a true "co-design" strategy for accelerated discovery, prediction, and control is an exciting grand challenge. Through co-design, the evolution of each capability will be guided by that of others so that observation, measurement, and simulation will be well matched to jointly tackle the study of complex phenomena. Achieving this qualitatively higher level of interdisciplinary cooperation and integration of major assets is essential if scientists are to meet the great national security and societal challenges of our time—health, energy, defense, and information management. DOE has a full spectrum of assets that could be used on these complex systems challenges, which include the next frontiers of extreme-scale computing. These opportunities are both significant and timely.

Over the next decade, the scientific community anticipates that if appropriate investments are made in research and development in computing technologies, increases in computing performance could be achieved by several orders of magnitude. These advances will have a major impact on scientists' ability to solve critical scientific and technological problems.

DOE has a long history of extending the frontiers of extreme computing applications and technology through its mission imperatives, which have included major national security application focuses (including nuclear weapons design and stewardship), major international user facilities, a variety of customers and markets, and the need for deep vertical (end-to-end system) integration in many missions. The National Nuclear Security Administration's (NNSA) need to continue on this path in extreme computing applications is well documented in the Advanced Simulation and Computing strategic progress towards resolving critical science and engineering issues, and moving to greater predictive systems capability in a non-underground test environment to ensure the safety, security, and effectiveness of the U.S. nuclear weapons stockpile. This breadth and depth has also resulted over the years in many significant demonstrations of leverage between nuclear security capability and broader national needs, including the Human Genome and the Global Climate Programs, high-fidelity combustion chamber modeling and design with the automotive industry, and impactful partnerships with the U.S. computing industry.

In this context, a technical workshop to discuss forefront questions in national security sciences and the role of high-performance computing was held on October 6-8, 2009, at the Hilton Washington, D.C. North/Gaithersburg hotel in Maryland. DOE's NNSA and the Office of Advanced Scientific Computing Research (ASCR) co-sponsored this collaborative workshop to identify leading scientific problems in national security science with a focus on problems requiring scientific computing capabilities at the extreme scale.

The purpose of this workshop was to address the following:

- Identify forefront scientific challenges in national security science and determine which of these could be aided by high-performance computing at the extreme scale.
- Establish how and why new high-performance computing capabilities could address issues at the frontiers of national security science within the DOE scope.
- Provide scientists the opportunity to influence the development of high-performance computing.
- Provide the national security science community with information on plans for the development of future high-performance computing capability by DOE.

The workshop was a forum for the exchange of ideas by scientists representing multiple disciplines—application scientists, computer scientists, and applied mathematicians—from the DOE Office of Science and other principal federal agency offices, NNSA laboratories, and several universities. One hundred and forty workshop participants provided the interdisciplinary expertise required to identify and address challenges in national security science and high-performance computing, with an emphasis on the use of extreme-scale computing for national security science research, advances, and discoveries. Also present as observers were several DOE Headquarters program managers.

Based on input from a multi-DOE laboratory organizing committee, technical panel discussions focused on six topical areas where extreme-scale computing is required to accomplish national security mission objectives. These panel areas were as follows:

- Multiphysics Simulation Problems
- Nuclear Physics
- Materials Science
- Chemistry
- Science of Nonproliferation
- Uncertainty Quantification and Error Analysis.

The principal common goal of these six technical panel discussions was to define significant future scientific challenges that require the timely development of extreme-scale computing. These scientific challenges were outlined in the introductory presentations at the workshop:

- *A Fascinating Time: Challenges and Opportunities* - Alan Bishop, Los Alamos National Laboratory
- *The Role of Computing in the Defense Program's "Predictive Capability Framework* - Robert Webster, National Nuclear Security Administration
- *Transformational Opportunities in Science* - James B. Roberto, Oak Ridge National Laboratory

- *The Imperative of High-Performance Computing: Pushing Back the Frontiers of Science and Technology* - William F. Brinkman, DOE Office of Science

- *"Co-Design" Defines a Future with Maximum Impact* - Paul Messina, Argonne National Laboratory

- *The Department of Energy and the Obama Administration: A Strategic Perspective* - Victor Reis, DOE Office of the Secretary.

Additionally, the co-chairs of the current DOE Exascale Initiative Steering Committee, Rick Stevens (Argonne National Laboratory) and Andrew White (Los Alamos National Laboratory), presented the current thinking on the computing technology roadmap and its effect on extreme-scale computing environments, and made suggestions on how to frame applications' requirements for exascale computing.

Workshop participants identified approximately 25 scientific challenges and associated priority research directions (PRDs) that span the full range of the workshop's scope:

Panel 1: Multiphysics Simulation Problems

1. **The path toward zero.** Work aimed at reaching the asymptotic goal of total nuclear disarmament; e.g., to reach the state of zero weapons in our nation's stockpile.

2. **Energy security for our nation.** Develop a path forward for assuring that sufficient energy supplies are available in the face of a climate-constrained world.

3. **Human/physical systems interaction.** Increase the levels of interaction between technologically sophisticated systems and environments, and human beings, which is one of the most challenging areas in multiphysics modeling.

4. **Critical resource protection.** Increase the focus on both the means by which the voluminous data gathered by remotely operated arrays of sensor networks can be captured and effectively analyzed and used, and on risk assessments of the vulnerability of these networks to either local- or system-wide failure, motivated by the U.S. dependence on those sensor systems.

5. **Enhanced industrial competitiveness.** Exploit the United States' continued preeminence in computing and simulations/modeling to regain its manufacturing preeminence, which depends critically on the willingness to engage in the challenges of solving critical simulation problems (and deploying their solutions) in areas that engage industrial competitiveness, many of which require multiphysics modeling.

Panel 2: Nuclear Physics

1. **Physics of fusion**. Compute the interactions and properties of nuclei that determine the energy-generating nuclear fusion reactions, as well as diagnostics in fusion environments such as those found in inertial confined fusion (e.g., at the National Ignition Facility), nuclear weapons, and astrophysical phenomena.

2. **Microscopic description of fission**. Compute pre-scission and post-scission quantities, and provide a consistent picture of the many fission data (cross sections, fission product properties, prompt fission neutrons and gamma rays, etc.).

3. **Theory of neutron reactions**. Provide an accurate, predictive capability of neutron cross sections of reactions on actinides and fission products from first-principles calculations.

4. **Simulation of integrated nuclear systems**. Provide necessary physics realism to other applications—such as nuclear reactors—that require the same types of nuclear physics input as the NNSA applications.

Panel 3: Materials Science

1. **Interfacial chemomechanics**. Gain a mechanistic understanding of environmental degradation caused by chemistry and stress through a combination of electronic structure, atomistic, and mesoscale simulations, and validation experiments.

2. **Atomistic simulation on engineering time scales**. Develop accurate multiscale models for microstructural evolution to meet certification challenges for deformation and failure under extreme conditions of radiation, mechanical loading, or environmental attack.

3. **Predict constitutive representation in multiphase materials**. Provide accurate representation of constitutive behavior of real materials under dynamic loading, including mechanisms of dissipation and damage in systems with realistic representations of materials heterogeneity.

4. **Electronic structure theory and simulation methods for nonadiabatic and strongly correlated systems**. Predictively treat strong electron correlation and electronic excitation with quantum simulation codes for ambient and extreme conditions.

Panel 4: Chemistry

1. **Advance the speed and accuracy of quantum chemistry modeling**. Develop goals for increasing accuracy to 1 kcal/mole where possible, capturing excited-state chemistry, and developing new algorithms that perform well on extreme-scale platforms.

2. **Enable predictive actinide chemistry**. Develop more accurate density functional theory functionals and equation-of-state calculations for f-electron materials.

3. **Reactive molecular dynamics with quantified accuracy**. Address the classical atomistic regime with the goal of enabling reactive molecular dynamics with quantified accuracy so that dynamics can be modeled at the micron/microsecond scale with a chemical accuracy of 1 kcal/mole.

4. **Seamless multiscale chemical reactivity**. Bridge length and time scales to enable seamless multiscale chemical reactivity modeling for high explosives. The seamless aspect can be achieved for the first time by running large-scale calculations with one method that directly overlaps the scale at which a coarser-grained method is used.

Panel 5: Science of Nonproliferation

1. **Optimization, inversion, and design**. Describe the fundamental mathematical and computer science challenges associated with how to better build, design, and interpret data from huge numbers of sensors.

2. **Proliferation process modeling**. Encompass the discipline of building sophisticated models of complex processes, including nuclear material production and weapons and materials trafficking.

3. **Information extraction and aggregatio**n. Address the especially difficult problem for the data associated with nonproliferation, as it consists of massive sets of streaming data, image data, spectral data, text data, and human observations.

4. **Information exploration**. Address the scientific challenge due to the nonphysical nature of the problem that places the analyst in a critical role.

5. **Statistics and machine learning for detecting rare and anomalous behavior**. Develop the technologies necessary for the computer to assist the analyst in identifying key findings in the data that would indicate anomalies representing suspicious behavior.

Panel 6: Uncertainty Quantification and Error Analysis

1. **Countering the curse of dimensionality**. Conduct the efficient and accurate exploration and exploitation of structure in high-dimensional, topologically complex spaces resulting from numerous uncertainties in physical model parameters, algorithmic approximations, databases, and observables.

2. **Intrusive/embedded uncertainty quantification**. Seek an alternative to exploring huge ensembles of simulations by altering the computational model to better align with uncertainty qualification needs.

3. **Uncertainty quantification in data-rich environments**. Develop methods to process high volumes of data without downgraded performance.

4. **Foundations in uncertainty qualification**. Focus on aspects such as extreme value theory, representing uncertainty, and providing a framework for assessing different uncertainty qualification methodologies.

5. **Combining disparate models and data sources**. Develop uncertainty qualification methods to combine simulation models at the quantum, meso, and bulk scales to infer material properties.

6. **Uncertainty quantification for emergency response**. Conduct time-constrained uncertainty qualification assessments using precomputed ensembles, reduced models, expert judgment, bounding methods, and/or other approaches.

Crosscutting Themes

In spite of the diversity inherent in the workshop, various crosscutting themes emerged:

- A new scale of technical challenge is posed by the need to address interdisciplinary complex systems (in scientific disciplines, and in urgent missions such as coupled energy-climate infrastructure, nonproliferation, etc.) to make improved prediction, control, and design tools available to decision-makers.

- Extreme-scale computing will continue to need to serve both huge computations at science frontiers (higher resolution, three-dimensional) and methods and models for enhanced predictive accuracy of both component processes and integral systems.

- Developing methodologies and algorithms capable of bridging length and time scales that span many orders of magnitude is an essential requirement for many applications anticipated at the exascale.

- The importance for all applications of advancing and using rigorous tools for uncertainty quantification, error analysis, optimization, inverse analysis, etc.

- The need to further develop and incorporate information science and technology methods—e.g., image analysis, machine learning, intelligent coarse-graining, stochastic modeling, model/hypothesis discovery and training—to process huge volumes of data, both observational and produced by extreme-scale simulations.

- The need for computational algorithms/codes for relevant physical phenomena, optimized to extreme-scale computing architectures.

- The importance of creating and enhancing mechanisms for deliberate teaming of expertise (application domains, computer science, applied math, etc.) to accelerate agile algorithm and code design for future extreme-scale computing architectures.

- Visionary partnerships (between universities, DOE laboratories, and industry) will be essential to achieve the envisioned extreme-scale computing impacts on national security science and mission frontiers, as well as to train a future workforce with the necessary interdisciplinary skills.

- Several areas of common interest to DOE's Office of Science and NNSA (e.g., extreme computing, climate-energy systems, materials under extreme conditions) were identified where progress should have major spin-off benefits to the U.S. industry.

Workshop participants enthusiastically and unanimously concluded that DOE's mission objectives require computing capabilities far beyond those currently available, and that an extreme-scale computing initiative—implemented within a deliberate "co-design" strategy—would be an excellent mechanism for accelerating the development of such capabilities.

TABLE OF CONTENTS

INTRODUCTION

Many scientific and engineering challenges would be enabled by the emergence of extreme-scale computing capabilities. The current projections of computer hardware component technologies in the next decade lead to the conclusion that extreme-scale computers (i.e., systems that are several orders of magnitude more powerful than today's petascale computers) will have different macro and micro architectures. These inevitable radical transitions anticipated in the immediate future for high-performance computing (HPC) architectures on the path to exascale and beyond (see "Architectures and Technology for Extreme Scale Computing"[1]) will be technologically disruptive for both hardware and software. The different computer architectures that will emerge will require substantial enhancements in mathematical models, numerical algorithms, programming models and languages, and system software.

This will certainly mark a generational technology change for HPC compared to the approach possible 15 years ago—namely, improving clock-speed to improve the time to solution, and significantly increasing the number of processors. The path to exascale will necessarily focus on very smart nodes, high parallelism, and effective systems software and programming mode. However, this challenge comes at a time when complex national security and societal challenges also require a far higher level of purposeful integration of major national assets. The science community is fortunate that the advanced computing growth of the last generation has reached a level of maturity such that it can now be fully incorporated into the "scientific method" for major new scientific and engineering frontiers, for remarkably complex systems, and with realistic goals of far greater levels of fidelity and prediction. Thus, this important time of challenging transition is also exciting for its huge potential in the science, technology, and engineering (ST&E) community, and impacts on societal health, security, and industrial competiveness. This will, however, require a dedicated path of investment and integration in terms of national facility assets—and equally important, of how the community works together in a strategy of "co-design" of computer science, algorithms and codes, architectures, *and* experimental and data acquisition design. These are major organizational, training, and management challenges and opportunities. To have the impact expected of us by the country on these urgent national security challenges, we must start now.

Significant and coordinated investments are needed to realize much of the research proposed in this report. These include large investments in experimental investigations, detection devices, and data acquisition (e.g., sensors). Likewise, the analysis and interpretation of results, and the development of a conceptual framework to guide understanding, prediction, and new investments in facilities depend significantly on focused investments, particularly in large-scale computing facilities, and the supporting infrastructure for data visualization, storage, retrieval, and transmission.

As described in the panel reports within this report, the scale of computing required for the necessary forefront research far exceeds currently available resources. Extrapolating the computer power needed to eliminate uncontrolled approximations and quantify controlled approximations and uncertainties, suggests for many national security (and associated scientific) frontiers that a major shift in capability will occur with computing capability and capacity several orders of magnitude beyond that available today.

[1] "Architectures and Technology for Extreme Scale Computing," December 8–10, 2009, San Diego, California. Workshop sponsored by DOE's Office of Advanced Scientific Computing and DOE/NNSA Office of Advanced Simulation and Computing.

Details of this workshop's overall charge from the U.S. Department of Energy (DOE Office of Advanced Scientific Computing Research [ASCR] and National Nuclear Security Administration [NNSA]), and the plenary presentations are provided on the extreme-scale computing workshop website at http://extremecomputing.labworks.org/nationalsecurity/index.stm. This workshop, "Scientific Grand Challenges in National Security: the Role of Computing at the Extreme Scale," was the final applications-focused workshop in the Scientific Grand Challenges series.[2] The workshop was distinct in that it deliberately brought together participants and application frontiers from DOE Office of Science and NNSA perspectives. Based on the workshop, there is clearly a great deal of intellectual and technical overlap between the communities, although historically they have often focused on somewhat different customer bases (user facilities and integrated codes for specific missions). In fact, this overlap has been exploited in the past to excellent effect (e.g., global climate modeling and simulation, the Human Genome Project, and the Advanced Combustion Engine R&D subprogram[3]), and the same leveraging is even more pressing for a practical decadal exascale initiative—for both frontier science and national security missions.

The purpose of this workshop was to address the following objectives:

- Identify forefront scientific challenges in national security science and determine which of these could be aided by HPC at the extreme scale.
- Establish how and why new HPC capabilities could address issues at the frontiers of national security science within the DOE scope.
- Provide scientists the opportunity to influence the development of HPC.
- Provide the national security science community with plans for the development of future HPC capability by DOE.

The workshop provided a forum for the exchange of ideas by scientists representing multiple disciplines—application scientists, computer scientists, and applied mathematicians—from the DOE Office of Science and other principal federal agency offices, NNSA laboratories, and several universities. Workshop participants (140 in attendance) provided the interdisciplinary expertise required to identify and address challenges in national security science and HPC, with an emphasis on the use of extreme-scale computing for national security science research, advances, and discoveries. Also present as observers were several DOE Headquarters program managers.

This workshop report is one of a series resulting from the Scientific Grand Challenges Workshops hosted by ASCR in partnership with other Office of Science programs. The workshop series focuses on the grand challenges of specific scientific domains and the role of extreme-scale computing in addressing those challenges. Dr. Paul Messina, Director of Science at the Argonne Leadership Computing Facility, is overseeing the workshop series.

---

[2] "Scientific Grand Challenges in National Security: the Role of Computing at the Extreme Scale," October 6–8, 2009, Washington D.C. Workshop sponsored by DOE's Office of Advanced Scientific Computing and the National Nuclear Security Administration.

[3] Formal title is the "Advanced Combustion Engine R&D subprogram of the U.S. Department of Energy's Vehicle Technologies Program."

Based on input from a multilaboratory organizing committee, technical panel discussions focused on six topical areas where extreme-scale computing is required to accomplish national security mission objectives. The topical areas were as follows:

- Multiphysics Simulation Problems
- Nuclear Physics
- Materials Science
- Chemistry
- Science of Nonproliferation
- Uncertainty Quantification and Error Analysis.

The principal common goal of these six technical panel discussions was to define significant future scientific challenges that require extreme-scale computing. These scientific challenges were outlined in the plenary presentations at the beginning of the workshop, and also placed in the context of the current DOE (ASCR and Advanced Simulation and Computing) perspective on the computing technology roadmap and its effect on extreme-scale computing environments. All panel members focused first on the key scientific questions in their areas, as well as the more immediate plans and milestones. Forefront problems that would benefit from the availability of extreme-computing applications over the next decade were then examined and research and computational requirements were estimated. The priority research directions (PRDs) identified by these panels follow.

Priority Research Directions

The six panels identified several grand challenges in science and engineering that—if met—would benefit national security missions and whose pursuit would be enabled by extreme-scale computing environments. The panels proceeded to also identify the following PRDs that must be pursued to tackle these grand challenges.

Panel 1: Multiphysics Simulation Problems

Multiphysics simulation is encountered in all missions supported by DOE. Multiphysics numerical simulation is not just a complex simulation—in its most simple form, it is modeling that involves two or more physical processes or phenomena that are coupled and often require disparate methods of solution. HPC has already dramatically expanded the accuracy and scope of processes that can now be modeled. Exascale computing has the potential to address new mission requirements for multiphysics simulation while presenting formidable challenges to physics models, numerical algorithms, and programming paradigms.

Over the past few decades, a great number of similar—as well as far more complex—multiphysics problems have been studied. While substantial progress has occurred in most instances, it is important to realize this area of computational modeling remains extremely challenging, with issues that are still largely unresolved. Examples of issues that remain unresolved include the following:

- Problems in which coupled physical processes have inherently different spatial and/or temporal attributes, leading to possibly conflicting discretizations of space and/or time.

- Problems in which the solution spaces for coupled physical processes are inherently distinct. Fluid dynamics in the presence of coupled radiation transport illustrate this issue: radiation transport is typically advanced in time separately from the fluid advection, and the underlying discretization of configuration space may look different for optimal solution of the fluid and radiation transport problems.

- Problems in which experimental data for validation of the computations may not be available.

In addition to the difficulties in developing accurate algorithms to address these issues, multiphysics simulations also pose serious challenges for performing computer simulations—a well-balanced computer architecture is needed.

One of the grand challenges facing scientific computation is the understanding and propagation of uncertainties in simulations. These challenges are particularly complex in multiphysics simulations, where scientists must distinguish between inherent uncertainties and physical processes omitted from the simulation.

This panel identified the following five key PRDs within the scope of multiphysics modeling:

1. **The path toward zero.** The effort to reach the asymptotic goal of total nuclear disarmament; e.g., to reach the state of zero weapons in our nation's stockpile. This goal entails not only significant political and sociological challenges, but also a number of technical challenges on a safe and secure path to the zero goal.

2. **Energy security for our nation.** Develop a path forward for assuring that sufficient energy supplies are available in the face of a climate-constrained world. This PRD involves several technical challenges ranging from the issues related to novel energy technologies to issues related to making existing energy technologies more economical, effective, and safe. Additional issues are those related to the verification of international agreements regarding the emission (and possible sequestration) of carbon dioxide (CO_2) and other greenhouse gases.

3. **Human/physical systems interaction.** This is one of the most challenging areas in multiphysics modeling. This PRD involves the increasing levels of interaction between technologically sophisticated systems and environments, and human beings. The challenges are not only prediction (and control) of the technology, but also prediction (and perhaps guidance) of human responses, on a variety of levels of interaction, time scales, and spatial scales.

4. **Critical resource protection**. The United States increasingly depends upon arrays of remotely operated sensor systems to gather information for a wide variety of critical issues ranging from national defense and national security to weather prediction, tsunami detection, and damage prediction and assessment. This dependence leads to increasing focus on both the means by which the voluminous data gathered by these sensor networks can be captured and effectively analyzed and used, and on risk assessments of the vulnerability of these networks to either local or system-wide failure.

5. **Enhanced industrial competitiveness**. Continued preeminence in computing and simulations/modeling offers a good opportunity for the United States to regain its manufacturing preeminence, which depends critically on our willingness to engage in the challenges of solving critical simulation problems (and deploying their solutions) in areas that engage industrial competitiveness, many of which require multiphysics modeling.

Panel 2: Nuclear Physics

Nuclear science is at the center of the NNSA program. The energy produced by nuclear processes is vital to the NNSA mission. Nuclear reactions are the critical source of energy in many applications, including National Ignition Facility (NIF) capsules, energy production and weapons, and in global threat reduction. Nuclear reactions can also be crucial in understanding and diagnosing the complex, high-energy environments present in all of these applications that are integral to controlling the reactions.

These nuclear processes are complex, many-body, strongly coupled quantum problems. Modeling and simulation of nuclear reactions and their role in applications, tightly coupled with experiments, have always played a key role in NNSA's mission. The science input to NNSA's program applications is heavily reliant on experiments combined with extrapolations and physical models marginally suitable to provide a starting point for extensive engineering that generates a body of empirical information. However, this body of information lacks the basic science underpinnings to provide reliable extrapolations beyond the domain in which it was produced. Furthermore, performing additional engineering tests, especially the tests that produced data in the extreme environments that uniquely characterize these applications, is no longer possible. The end of nuclear weapons testing has required improvements to the predictive capabilities of codes simulating the applications for well-known, well-characterized cases as well as incompletely known cases.

Developments in HPC, computational physics, applied mathematics, and nuclear theory have combined to make dramatic advances in the theory of fission, fusion, and nuclear reactions. Thermonuclear fusion reactions involve light nuclei, where clustering and correlations play a critical role. For these reactions, it is just now becoming possible to use *ab initio* (i.e., from first principles) approaches, solving the quantum few- and many-body scattering problem with a realistic inter-nucleon force. Nuclear structure, decays, and reaction phenomena are then predicted from these solutions. Predictive simulations of key reactions would be extremely valuable in designing experiments to diagnose; for example, hydrodynamic instabilities at the NIF at Lawrence Livermore National Laboratory. These simulations would also be of great value in understanding related explosive astrophysical environments.

For the heavier nuclei undergoing fission and related processes, the nuclear mean field is well developed and describes a valid starting point for fission calculations. However, in contrast to many other quantum applications, the superfluid (pairing) correlations play a critical role here. Recently, large-scale simulations have started to provide a realistic description of reactions, as well as spontaneous and induced fission. In many cases, including reactions on fission products, only theoretical and computational approaches can advance researchers' understanding of these phenomena. Including important quantum fluctuations in the nuclear fission process, and expanding beyond the current mean-field approach is demonstrably an exascale problem.

These recent developments in theory and computation make it possible to project that calculations describing the details of nuclear reactions relevant to NNSA's program can be completed. They can be

performed with the highest physical realism, starting from the correct fundamental description of nuclear physics—thereby providing true predictive capability with regards to the nuclear phenomena. Exascale computing will make these claims realizable.

Based on this progress, the panel identified the following four major PRDs that will be achieved in nuclear science using extreme-scale computing over the next decade:

1. The **physics of fusion**, which includes computing the interactions and properties of nuclei that determine the energy-generating nuclear fusion reactions, as well as diagnostics in fusion environments such as those found in inertial confined fusion (for example, at the NIF), nuclear weapons, and astrophysical phenomena.

2. The **microscopic description of fission**, which includes computing pre-scission and post-scission quantities, and providing a consistent picture of the many fission data (cross sections, fission product properties, prompt fission neutrons and gamma rays, etc.). These processes are critical for energy, global security, and weapons applications.

3. The **theory of neutron reactions** that will provide an accurate, predictive capability **of neutron cross sections of reactions on actinides and fission products** from first-principles calculations.

4. The **simulation of integrated nuclear systems** to provide necessary physics realism to other applications, such as nuclear reactors, that require the same types of nuclear physics input as the NNSA applications. High-fidelity simulations of these problems must integrate several phenomena, including accurate particle transport, reactions with "complete" correlations and uncertainties, and nuclear burn. All areas require dramatic improvements of current computational capabilities. Researchers anticipate that simulations of these integrated systems, for either fission or fusion reactors, will be a necessary component of exploring possible future sources of energy.

Panel 3: Materials Science

In the last decades, the process of certifying the performance of nuclear weapons has changed from integrated tests to science-based certification relying on large-scale simulations appropriately validated with experimental data. While significant advances have been made in incorporating fundamental materials physics into integrated codes, uncertainties in addressing materials properties, especially at the mesoscale, remain key in eliminating researchers' current empirical knobs. National-security mission requirements demand that materials perform predictably in extreme environments (high pressure, high strain rate, and hostile irradiation and chemical conditions). In particular, dynamic materials properties are critical to understanding weapons performance—and materials compatibility and aging phenomena represent significant challenges.

The advent of exascale computing provides an opportunity to advance the frontiers of scientific understanding of materials performance, including the effects of materials fabrication and processing on ultimate performance and materials failure (i.e., process-aware performance). The possibility now exists to achieve the predictive capabilities required to manipulate microstructure and interfaces at the grain scale. This will enable the design and development of extreme-environment tolerant advanced materials. The exascale challenge is direct access to the mesoscale frontier, predicting mechanisms-governing materials properties in micron-scale volumes over milliseconds. These NNSA-specific issues are a subset

of broader materials challenges, especially in extremes (e.g., materials needs for nuclear energy). Success in this endeavor will require strong coupling and feedback among experiment, theory, and computation.

The panel identified the following four PRDs in which exascale computing resources could have a specific and transformational impact:

1. **Interfacial chemo-mechanics**. Gain a mechanistic understanding of environmental degradation coupling chemistry and stress through a combination of electronic structure, atomistic, and mesoscale simulations, and validation experiments. Examples of such phenomena include stress-corrosion cracking, mechanochemistry of foams, nanostructured materials, and protective coatings. With exascale computers, researchers anticipate direct comparability could be achieved with science-driven experiments—microstructure, millisecond resolution could be gained and materials performance could be tailored for specific environments.

2. **Atomistic simulation on engineering time scales**. Accurate multiscale models for microstructural evolution are needed to meet certification challenges for deformation and failure under extreme conditions of radiation, mechanical loading, or environmental attack.

3. **Predict constitutive representation in multiphase materials**. Provide accurate representation of constitutive behavior of real materials under dynamic loading, including mechanisms of dissipation and damage in systems with realistic representations of materials heterogeneity.

4. **Electronic structure theory and simulation methods for nonadiabatic and strongly correlated systems**. Predictively treat strong electron correlation and electronic excitation with quantum simulation codes for ambient and extreme conditions.

Common to all of these PRDs is the need to understand fundamental mechanisms and their consequences across multiple spatial and temporal scales focused on the scale of grains and interfaces. Strong synergies with the Chemistry panel were identified in two areas: the need to bridge atomic scale spatial and temporal phenomena to more integrated scales, and the need to understand excited state phenomena and dynamics, beyond ground-state calculations.

Panel 4: Chemistry

The breaking and forming of bonds between atoms has an important role in NNSA's general interest in material properties in extreme environments. In computational models, this occurs at the quantum level where electrons are included explicitly, at the atomic level in molecular dynamics and Monte Carlo methods that include bond-order effects, and in mesoscale and even continuum models that include chemical reactions in their Hamiltonians or via rate equations. Increasing the accuracy and predictive power of such models on next-generation extreme-scale computers presents several new opportunities.

The Chemistry panel identified three basic science challenges of critical interest to NNSA that rely on accurate chemistry modeling. The first is to understand the response of target materials in inertial confinement fusion devices, such as NIF and the Z-pinch machine. The second is to assess the safety, control, and performance of high explosives used in nuclear weapons. The third is to predict the aging of weapon components. Additionally, the chemistry panel identified two basic science challenges of broader

national security interest. The first is to design molecular sensors via first-principles modeling. The second is to predict how actinides and their associated compounds will behave in environmental settings.

Based on these science drivers, the Chemistry panel selected four crosscutting PRDs that could be enabled by the availability of computing resources at the extreme scale, and thus have a significant impact on all five of the science challenges:

1. **Advance the speed and accuracy of quantum chemistry modeling**, with goals for increasing accuracy to 1 kcal/mole where possible, capturing excited-state chemistry, and developing new algorithms that perform well on extreme-scale platforms.

2. **Enable predictive actinide chemistry**, which will require more accurate density functional theory functionals and equation-of-state calculations for f-electron materials.

3. Address the classical atomistic regime with the goal of **enabling reactive molecular dynamics with quantified accuracy** so that the dynamics can be modeled at the micron/microsecond scale with chemical accuracy of 1 kcal/mole. This will require automated methods of generating reactive force fields for a broad range of materials by incorporating quantum-calculated information and training reactive models appropriately.

4. Bridge length and time scales to **enable seamless multiscale chemical reactivity** modeling for high explosives. The seamless aspect can be achieved for the first time by running large-scale calculations with one method that directly overlaps the scale at which a coarser-grained method is used. This will enable direct verification of the coupling algorithms used between methods and will thus create a direct predictive modeling capability for high explosives.

Panel 5: Science of Nonproliferation

DOE's nuclear nonproliferation effort is focused on three primary areas: detecting nuclear materials and programs; securing nuclear weapons and material around the world; and disposing of surplus nuclear and radiological materials. While the problem of securing weapons is very operational and the disposal problem is already heavily studied, the problem of detection is an area that has tremendous potential to be improved by application of HPC. Current detection methodologies rely heavily on the analyst to compile large amounts of data from different sources to make high-consequence decisions. HPC, especially at the exascale, could impact this area very broadly throughout the analysis process by increasing the data-gathering ability of sensors, to information aggregation and storage, and finally to sophisticated analysis techniques that allow analysts to make sophisticated queries of the data. In the mission space of DOE, these technical challenges are unique because they deal less with well-defined physical models involving differential equations and more with the subtle problems of working with abstract models and data types.

Panel members found the questions related to the complex data queries above led to five specific PRDs:

1. **Optimization, inversion, and design** describes the fundamental mathematical and computer science challenges associated with how to better build, design, and interpret the data from sensors. Possessing higher-quality data from the nonproliferation community's sensors and a more fundamental understanding of the uncertainties associated with data derived from these sensors would be a tremendous advantage to the nonproliferation community.

2. **Proliferation process modeling** encompasses the discipline of building sophisticated models of complex processes, including nuclear material production and weapons/materials trafficking.

3. **Information extraction and aggregation** is an especially difficult problem for the data associated with nonproliferation because it consists of massive sets of streaming data, image data, spectral data, text data, and human observations. Bringing these and other disparate data types together is an especially challenging problem.

4. The broad category of **information exploration** is the fourth PRD. Again, the science of nonproliferation distinguishes itself from many of the other areas of expertise due to the nonphysical nature of the problem that places the analyst in a critical role: an individual's ability to interact with diverse types of media to enable key decisions is critical to the analysis process.

5. **Statistics and machine learning for detecting rare and anomalous behavior** includes all of the technologies that are necessary for the computer to assist the analyst in identifying key findings in the data that would indicate anomalies representing suspicious behavior.

Panel 6: Uncertainty Quantification and Error Analysis

The field of uncertainty quantification describes the reliability of scientific inferences. Uncertainty qualification seeks to augment the best available answer with an understanding of the range of plausible bounds to that answer—and their relative likelihoods—resulting from a given action or inaction. In the context of national security, inferences are typically made about particular physical, biological, chemical, or social systems with the aid of some form of computational model. These computational models may be derived from first principles (e.g., quantum molecular dynamics), mathematical-physical descriptions (e.g., global circulation, nuclear weapons, inertial confinement fusion, etc.), empirical relationships (such as agent-based models), or some combination. Uncertainty qualification provides the integrating framework, theory, and methodology for making inferences about such systems using computational models, theoretical approximation, experimental and/or observational data, and expert judgment. Thus, uncertainty qualification produces uncertainty estimates regarding important predictions, as well as a framework for assessing the impact of various actions against the system of interest. For example, uncertainty qualification methods can determine what additional observational data will give the largest reduction in prediction uncertainty, will indicate what parameters or approximations in a physical model drive the overall uncertainty in system prediction, or will quantitatively describe how a potential mitigation strategy may affect global CO_2 concentrations.

HPC and computational modeling play a dominant role in shaping the methodological developments and research in uncertainty qualification. Depending on the complexity of the uncertainty qualification investigation, anywhere from 10^2 to 10^8 runs of the computational model may be required. Thus, uncertainty qualification investigations may require extreme-computing environments (e.g., exascale) to obtain results in a useful time frame, even if a single run of the computational model does not require such resources.

With these computational considerations in mind, panel members identified the following six PRDs:

1. **Countering the curse of dimensionality**, which involves the efficient and accurate exploration and exploitation of structure in high-dimensional, topologically complex spaces resulting from numerous

uncertainties in physical model parameters, algorithmic approximations, databases and observables. Such research will allow more efficient use of large ensembles of simulations, and will improve both the exploration and the modeling of computational model response to changes in inputs.

2. **Intrusive/embedded uncertainty quantification** seeks an alternative to exploring huge ensembles of simulations by altering the computational model to better align with uncertainty qualification needs. For example, computational models could be constructed with adjoints to produce derivative information. Such approaches can greatly reduce the number of simulations required for uncertainty qualification.

3. **Uncertainty quantification in data-rich environments** that must process high volumes of data without downgraded performance. Here, the focus is on developing theory, methods, and algorithms for uncertainty qualification that can leverage new computing architectures to make efficient use of high-volume data from sensors, diagnostics, or simulation output.

4. **Foundations in uncertainty qualification**, which focuses on aspects such as extreme value theory, representing uncertainty and providing a framework for assessing different uncertainty qualification methodologies.

5. **Combining disparate models and data sources**. For example, uncertainty qualification methods are needed for combining simulation models at the quantum, meso, and bulk scales to infer material properties. A wide variety of models (circulation, ecological, hydrological, etc.) are needed to infer the response in CO_2 concentration to mitigations.

6. **Uncertainty quantification for emergency response**. In some applications, a rapid assessment of uncertainties is needed. Examples include emergency response to hurricanes or a terrorist attack. This PRD focuses on making time-constrained uncertainty qualification assessments, using pre-computed ensembles, reduced models, expert judgment, bounding methods, and/or other approaches.

Crosscutting Challenges

Although the focus of this workshop was on applications opportunities and challenges, the workshop participants also identified issues and challenges that result from the complexity and scale of the applications enabled by extreme-scale computing systems. Those challenges are primarily in the fields of mathematical models, numerical algorithms, software technology, and computer architectures. Additional challenges involve training and multidisciplinary and multi-institutional partnerships.

- A new scale of technical challenge is posed by the need to address interdisciplinary complex systems (in scientific disciplines and in urgent missions such as coupled energy-climate infrastructure, nonproliferation, etc.) to make improved prediction, control, and design tools available to decision-makers.

- Extreme-scale computing will continue to need to serve both huge computations at science frontiers (higher resolution, three-dimensional) and methods and models for enhanced predictive accuracy of both component processes and integral systems.

- Developing methodologies and algorithms capable of bridging length and time scales that span many orders of magnitude is an essential requirement for many applications anticipated at the exascale.

- The importance for all applications of advancing and using rigorous uncertainty quantification, error analysis, optimization, inverse tools, etc.

- The need to further develop and incorporate information science and technology methods—e.g., image analysis, machine learning, intelligent coarse-graining, stochastic modeling, model/hypothesis discovery and training—to deal with the huge volumes of data, both observational and output of extreme-scale simulations.

- The need for computational algorithms/codes for relevant physical phenomena, optimized to extreme-scale computing architectures.

- The importance of creating and enhancing mechanisms for deliberate teaming of expertise (application domains, computer science, applied math, etc.) to accelerate agile algorithm and code design for future extreme-scale computing architectures.

- Visionary partnerships (between universities, DOE laboratories, and industry) will be essential to achieve the envisioned extreme-scale computing impacts on national security science and mission frontiers, as well as to train a future workforce with the necessary interdisciplinary skills.

- Several areas of common interest to DOE's Office of Science and NNSA (e.g., extreme computing, climate-energy systems, materials under extreme conditions) were identified where progress should have major spin-off benefits to the U.S. industry.

PANEL REPORTS

MULTIPHYSICS MODELING

NUCLEAR PHYSICS

MATERIALS SCIENCE

CHEMISTRY

THE SCIENCE OF NONPROLIFERATION

UNCERTAINTY QUANTIFICATION AND ERROR ANALYSIS

MULTIPHYSICS MODELING

Co-Leads: Robert Rosner, University of Chicago
Michael Zika, Lawrence Livermore National Laboratory

Panel Members: Ricardo Betti, University of Rochester; Scott Delp, Stanford University; Edward Dendy, Los Alamos National Laboratory; Paul R. Drake, University of Michigan; Rao Kotamarthi, Argonne National Laboratory; Don Lamb, University of Chicago; Edison Liang, Rice University; Parvis Moin, Stanford University; Desmond Pilkington, Lawrence Livermore National Laboratory; Fred Streitz, Lawrence Livermore National Laboratory; John A. Turner, Oak Ridge National Laboratory; Robert Webster, Los Alamos National Laboratory; George Zimmerman, Lawrence Livermore National Laboratory; and Dennis Hewett, Lawrence Livermore National Laboratory

CURRENT STATUS

Multiphysics simulation is encountered in all missions supported by the U.S. Department of Energy (DOE). Multiphysics numerical simulation is not simply a "complex simulation"—in its simplest form, it is modeling that involves two or more physical processes or phenomena that are coupled, and often require disparate methods of solution. High-performance computing (HPC) has dramatically expanded the accuracy and scope of processes that can now be modeled. Exascale computing has the potential to access a new mission space for multiphysics simulations while presenting formidable challenges to physics models, numerical algorithms, and programming paradigms.

As experiments become more difficult or impossible (as in the case of the cessation of nuclear testing), multiphysics simulation has assumed an increasingly important role for understanding the most complex coupled physical processes. This increases the need to migrate from computational models built on empiricism and calibration to physics-based models built on a fundamental understanding of the underlying physical processes. Such "predictive capability" in multiphysics simulations opens the possibility of predicting performance of a multiphysics system *and* providing quantifiable error bounds for the performance of that system.

Developments in HPC, computational physics and mathematics, and fundamental theory have combined to enable significant advances in simulation capability for complex multiphysics application areas. At the same time, increasingly sophisticated experiments revealing underlying physical phenomena have also advanced fundamental theory. Combined, they have opened the possibility for predictive capability when modeling complex multiphysics processes. Current research synthesizes these advancements into multiphysics systems modeling capabilities to provide predictive capability in numerous application areas. Such a predictive capability can enable the transition from fundamental understanding to design and optimization of practical engineered systems *if* simulations are sufficiently accurate and enable a reasonable turnaround time. To apply these capabilities to areas relevant to national security drives the scientific community toward exascale computing.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

Over the past few decades, a great number of multiphysics problems have been studied, and while substantial progress has occurred in most instances, it is important to realize that this area of computational modeling remains extremely challenging, with issues that are still largely unresolved. Some examples of the issues include the following:

1. **Problems in which coupled physical processes have inherently different spatial and/or temporal attributes**, leading to possibly conflicting discretizations of space and/or time. A prototypical example is the combination of advection and diffusion in the Navier-Stokes equation, leading to both parabolic and hyperbolic terms. While advection is characterized by a finite propagation speed for information, diffusion inherently contains an infinite propagation speed.

2. **Problems in which the solution spaces for coupled physical processes are inherently distinct**. Fluid dynamics in the presence of coupled radiation transport illustrates this issue: fluid advection is typically dealt with by a deterministic continuum hyperbolic transport equation in real (configuration) space, while the radiation transport may have to be dealt with by solving a (coupled) Monte Carlo particle transport equation in a higher-dimensional solution space (consisting—in the most simple case—of the outer product space defined by three-dimensional configuration space and a one-dimensional energy space). The latter equation is typically advanced in time quite separately from the time integration of the continuum equation; and the underlying discretization of configuration space may well look quite different for optimal solution of the fluid and radiation transport problems.

3. **Problems in which experimental data for validation of the computations may not be available**. This problem is classically most prominent in radiation hydrodynamics, primarily because the conditions under which radiation and matter are strongly coupled are extremely difficult to replicate in terrestrial laboratories.

In addition to the difficulties in developing accurate algorithms to address these issues, multiphysics simulations pose serious challenges for performing computer simulations. Returning to the example of coupled radiation-hydrodynamics, the physical processes in the simulation impose inherently distinct demands on the computer architecture. Hydrodynamics are characterized by moderate floating-point computations with regular, structured communication. Monte Carlo particle transport is characterized by intense fixed-point computations with random communication. As a result, multiphysics simulations typically require a balanced computer architecture in terms of processor speed, memory size, memory bandwidth, and interconnect bandwidth, at a minimum.

As a result of discussions and consideration of recent progress, the panel identified four key priority research directions (PRDs) in multiphysics modeling that will drive exascale computing in the coming decade. The first PRD is the **path toward zero**, which includes a predictive capability for the physics understanding of nuclear weapon systems and the use of that capability to support complex-wide assessments and optimization. This PRD is named in recognition of President Obama's vision of "walking a path toward a world without nuclear weapons." The second PRD is **energy security**, which addresses the nation's need for secure and reliable energy production and distribution. The third PRD is **critical resource protection**, which identifies the need to understand and protect critical infrastructure—such as satellites—from disruption. To protect this infrastructure, scientists must first realize a fundamental understanding of the threats. The fourth PRD is **enhanced industrial competitiveness**,

which includes maintaining preeminence in engineering and manufacturing by deploying simulation capability in arenas critical to industrial competitiveness.

Each PRD is presented in the following sections, which describe the scientific and computational challenges involved in each area, summarizing the key research directions and expected scientific and computational outcomes, and describing the potential impact this research may have on multiphysics problems that arise both within the National Nuclear Security Administration (NNSA)/Advanced Simulation and Computing (ASC) environment and in related spheres of interest.

Path Toward Zero

The "path to zero" refers to the effort espoused by every administration since the advent of nuclear weapons to reach the asymptotic goal of total nuclear disarmament; i.e., to reach the state of zero weapons both nationally and globally. In his State of the Union Address on January 27, 2010, President Obama stated:

> I've embraced the vision of John F. Kennedy and Ronald Reagan through a strategy that reverses the spread of these weapons and seeks a world without them. To reduce our stockpiles and launchers, while ensuring our deterrent, the United States and Russia are completing negotiations on the farthest-reaching arms control treaty in nearly two decades.

Work toward the goal of zero weapons entails not only significant political and sociological challenges, but also a number of technical challenges, without whose solution the path to zero will be impossible to reach.

Three high-level national security objectives have not been addressed, even with access to terascale and petascale computing:

1. Cradle-to-grave assessments to support complex-wide process and resource optimization

2. System-scale simulation support for national policy decisions on stockpile changes or reductions

3. System-scale physics understanding of nuclear weapons, including complex interacting microscale processes.

These objectives are made feasible when exascale computing is used to combine and dramatically extend results previously obtained at the terascale and petascale. Reaching these objectives would achieve a fundamental understanding that may eliminate the technical barriers present in realizing the path to zero.

The drive to terascale in the original ASC Program in the mid-1990s culminated in pioneering proof-of-principle simulations of nuclear weapons in three dimensions. At terascale, some traditional *ad hoc* physics models were replaced with improved physics-based models, reducing reliance on empiricism in modeling. While the computational meshes were adequate for entry-level three-dimensional simulations, memory and execution speed at terascale were only sufficient for initial attempts to address geometric complexity or to eliminate key *ad hoc* approximations. In 2004, the NNSA ASC Program noted this success:

> Given the current maturity of physics models of ASC codes, capability calculations today are measured within a window of 10 to 100 teraOPS….However, the major driver for capability machines of the future, petaOPS and

> beyond, will be the need to model the additional physical phenomena that describe the time evolution of a nuclear device at the level of detail and fidelity that sufficient confidence demands." (Kusnezov 2004).

The establishment of petascale computing over the last decade, now nearing fruition, is enabling the identification of the underlying physical processes to remaining *ad hoc* treatments that was previously only postulated. Increased mesh resolution is allowing more sophisticated multiphysics models to be applied, at least locally, within major system components. In at least one of these issues, petascale simulations in two dimensions and three dimensions are leading to new understanding of long-standing uncertainties. Petascale computing is thus identifying and setting simulation requirements for key physics issues.

Even at the petascale, the memory and execution speed are being found to be insufficient to treat highly nonlinear interactions between complex multiphysics processes that are intimately tied to localized structures and properties. System-scale simulations with sufficient resolution and physics fidelity to replace reliance on empiricism with physics-based models must await the advance to exascale computing.

Cradle-to-Grave Assessments for Complex-Wide Optimization

DOE's Accelerated Strategic Computing Initiative (now ASC) Program successfully delivered "button-to-boom" simulation capability enabled by terascale computing. The simulation capability needed by the Stockpile Stewardship Program4 to support complex-wide optimization must go beyond this focus to enable a broader context the scientific community refers to as "cradle-to-grave" (see Figure 1). This broader view includes the initial conditions, structure, and multiscale material properties of the systems, as they are established by the manufacturing processes and subsequently modified by environmental effects, handling, and storage. This view also extends to realistic assessments of the end-state of the system, whether it is safe and secure dismantlement, or used accidentally, maliciously, or intentionally. In general, exascale simulation capabilities are needed to provide validated predictions of the system safety, security, and effectiveness implications of as-built, as-assembled, as-aged, or as-disassembled systems, including multiscale structure and material properties, at any stage in the system's lifecycle.

The emphasis on microscale material structure and properties stems from the new understanding now being obtained at the petascale on how these small-scale conditions relate to key physics issues. This understanding forms the foundation for a manufacturing-aware simulation capability. Exascale system-level simulations build on techniques developed at the petascale to identify the dominant physical processes, and then model their interactions across multiple physical and time scales. The added emphasis on the end-of-life state includes calculation of the detailed outputs (e.g., radiation and debris) produced because these drive the impact on real-world environments following planned or unplanned system detonation. For intended use, this concept includes the military effectiveness of the system and any collateral consequences. For unplanned use, whether accidental or malicious, this concept includes assessing the extent and severity of damage, the effectiveness of protection systems or procedures, and, if necessary, inference of the explosion source and identity of the perpetrators of the event.

This objective requires and builds on exascale simulations of atomic-scale and microscale dynamic material properties, achieved at the petascale. In turn, this objective will produce the driving initial

[4] Bureau of Arms Control, U.S. Department of State, Washington, D.C.

conditions for exascale simulations of the response of defense and civilian structures and environments to explosive releases of nuclear energy.

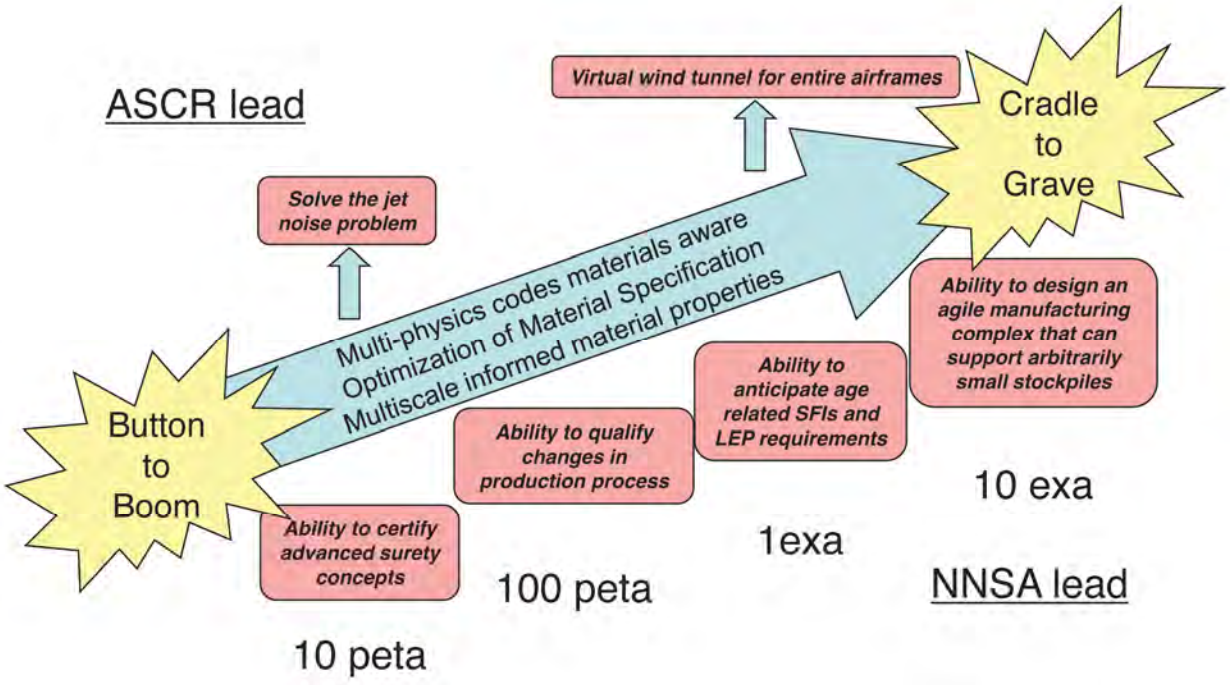

Figure 1. Reinvest in the "peace dividend" to enhance U.S. industrial competitiveness on the path to zero weapons. Image courtesy of Robert Webster (Lawrence Livermore National Laboratory).

System-Scale Simulation for National Policy Decisions

The Stockpile Stewardship Program greatly increased computational capability following the end of nuclear testing. This growth in computing power and weapons simulation capability through the terascale into the petascale has provided robust support for stockpile activities, including Life Extension Programs and exploration of potential refurbishment or replacement options for specific systems. The new challenge, requiring an advance in exascale computing, is to support the sustainable stockpile as national policy decisions are contemplated that may significantly reduce the numerical number of weapons or the diversity of the systems in the stockpile. The new requirement will be for high-confidence, system-scale simulations to assess the safety, security, and effectiveness of the remaining stockpile.

With the possibility of a much smaller stockpile, simulations supporting policy decisions will likely have to address these characteristics for smaller classes of systems or possibly individual systems. Exascale simulations will be needed in these circumstances to provide validated predictions for much more closely defined configurations. It may even be necessary to conduct suites of system-scale simulations for statistical studies of potential configuration changes resulting from aging, storing, and handling stockpiles in varying environmental conditions.

In parallel, system-scale simulations may be needed to explore the consequences of postulated systematic effects, such as the hypothetical degradation of some key component. The possibility of fewer weapons or weapon systems will make it more important to assess the impact of departures from as-designed

properties or configurations than when a much larger weapon count allowed for a broader sample of statistical variations.

The need for three-dimensional simulations of specific systems or configurations with subtle deviations from nominal will require exascale computing where these detailed initial conditions can be taken into account. To address the potential impact of specific deviations from nominal requires a simulation capability with enough precision—both in mesh resolution and physics fidelity—to identify any anomalous system performance.

System-Scale Physics Understanding

The advance from terascale to petascale enables the identification and modeling of new classes of physical processes, some of which were historically treated with *ad hoc* approximations. The successes scientists are beginning to see at the petascale level are establishing simulation requirements for these issues, both in terms of mesh resolution and computational physics models. Sometimes the mesh resolution requirements can be inferred from petascale simulations that start to resolve physical processes that could not be seen directly at lower resolution. In some cases, it has been possible to anticipate what the impact of significantly higher resolution would be in three dimensions by conducting highly refined calculations in two dimensions or by simulating in three dimensions a very small representative portion of a larger system component. The difficulty is that some essential phenomena are inherently three-dimensional and cannot be properly represented in two dimensions, or that a simulation of a localized region misses interactions with adjacent regions that fundamentally affect the overall system behavior. Even as these efforts are succeeding, they are still limited in that the relevant phenomena involve multiphysics interactions with microscale structure and material properties. Again, even at the petascale, a simulation for particularly complex phenomena interacting with a realistic configuration may only be able to treat a small spatial extent at the highest physics fidelity. This means there is only a very limited capability to treat complex nonlinear interactions across the system scale. The need for system-scale simulations at the highest resolution and physics fidelity lead directly to the requirement for exascale computing.

An illustrative example of the multiphysics drivers of exascale computing can be found in the need to represent complex nonlinear interactions in the explosion of a Type Ia supernova. Pioneering near-petascale simulations of such an explosion in three dimensions (Jordan et al. 2008) show the observed properties of such supernovae can be matched if the detonation of the star begins at a point away from the geometric center of the star. An off-center detonation is reasonable because there are chaotic turbulent flows in the cores of these stars at velocities comparable to laminar nuclear flame speeds. Even with this computing power, the simulations assumed idealized initial conditions and had to use an approximate treatment for subgrid-scale turbulence. The authors of this panel report note that “the nuclear flame front in the deflagration phase of Type Ia supernovae is … unresolvable in any whole-star simulation.” They also note that “the appropriate way to treat turbulent nuclear burning is an open question, and further studies are needed in order to answer definitively this question.” Removing these assumptions and treating more realistic initial conditions for Type Ia supernova explosions with sufficient resolution and physics fidelity would push the required computing resources far beyond the petascale.

While the scales and physical processes in weapons systems are very different from supernova explosions, the need to represent complex nonlinear interactions is similar. Scientists may be successful at the petascale in simulating the effect of complex multiphysics phenomena acting on a small portion of

a weapon system. This is a great advance over the previous generation because it allows modeling of previously unidentified physical processes. For some physical phenomena, this may be sufficient to advance knowledge of system safety, security, and effectiveness. However, for certain physical processes, it is necessary to account for complex nonlinear interactions across the entire system-scale configuration. Extending simulations that require petascale computing for a *portion* of the system to encompass the *entire* system results in the need for exascale computing. The need is particularly acute when detailed initial geometric configurations and microscale material properties must be treated. Futhermore, while validation of the simulation tools at the microscopic level is usually feasible (and indeed essential), the feasibility of experimental validation becomes increasingly more challenging as scientists move to the subsystem and ultimately to the full systems level. This situation is made particularly acute because such validation experiments remain the "gold-standard" for ensuring the correctness of simulation results. The challenge posed by this difficulty—one that has driven the invention of a number of clever "proxy" experiments for higher systems-level experimental validation—at both the intellectual and financial levels should not be underestimated.

Exascale computing of nuclear weapons is essential to extend the advances in simulation capability achieved at the terascale and petascale to meet national security objectives. This simulation capability is needed for cradle-to-grave safety, security, and effectiveness assessments at any stage of system life; high-confidence simulations to support national policy decisions about stockpile changes; and weapons physics simulations to extend advances in key physics issues by including complex nonlinear interactions driven by microscale effects across the system scale. This simulation capability supports three broad national security objectives: optimize complex-wide processes and resources; sustain a stockpile that may be reduced in size or diversity through national policy changes; and provide a high-confidence system-scale weapons physics simulation capability.

Energy Security

A complex set of issues emerges from the near certainty that the overall energy demands within the United States will continue to increase over time, and the requisite energy supplies in a world constrained by global climate concerns and by the political instability of many key energy suppliers may well be threatened. This complexity involves a mix of technology problems to be solved, from developing capable tools to verify compliance with international climate treaty obligations to improving the efficiency of energy use and distribution, and to the development of new carbon-neutral energy technologies. A broad range of physical science disciplines must be engaged in virtually all of these problem areas; the predictive power of modern high-performance modeling and simulation will be essential to their successful application.

Carbon Treaty Verification

A key element in any scenario for limiting the impact of global warming driven by modern human civilization is the establishment of verifiable goals for the reduction of carbon dioxide (CO_2) and other greenhouse gas emissions on a nation-by-nation basis (Loreti et al. 2001). For such goals to be verifiable, scientists need to establish credible baselines for current greenhouse gas emissions and enable the monitoring of future emissions. The technology challenges here involve the design and development of ground- and space-based sensor technologies capable of effective monitoring of the emissions and optimal deployment and safeguarding of these sensor networks. These networks must be capable of sensing the appropriate greenhouse gases in local environments and monitoring local weather conditions.

Ultimately, the verification system must then reduce and analyze extremely voluminous (petabyte and above) data from the field to provide insight to decision-makers.

To illustrate the research issues to be addressed, scientists focus on the coupling of real-time sensor networks to real-time regional and global weather and climate models. National treaty obligations naturally refer to the cumulative greenhouse emissions from treaty participants, implying that some sort of spatial accumulation over the territory of such participants must be conducted. Such accumulations, when derived from localized point source data, are clearly affected by weather conditions. For example, atmospheric mixing does not recognize national boundaries, and thus scientists face the challenge of mastering the combined modeling of local atmospheric transport with transport on scales of global weather and climate at levels of accuracy sufficient to meet the demands of climate treaty obligations.

The desired outcome of this effort would result in the design and production of ground- and/or space-based sensor networks capable of providing both weather and greenhouse gas concentration data. That sensor network data would then be integrated with a new generation of weather and climate computational models designed specifically to derive spatially and temporally accumulated greenhouse gas emission levels on the national level. U.S. national foreign policy—and national security—is conditioned on the ability to verify commitments made by other nations in treaties with the United States. In the absence of such capabilities, an enforceable "carbon treaty" is an oxymoron; in the presence of such capabilities, United States officials would instead have the tools for managing national and other countries' treaty obligations.

Nuclear Fission

If nuclear fission is to remain (or grow) as a key element in the nation's energy portfolio, scientists need to address key elements relevant to nuclear energy: the capital costs and the safety and security of fuel enrichment and fabrication, nuclear reactor facilities, and spent fuel disposition. In areas of safety and economics of nuclear reactors, advancements are needed in the fundamental understanding of material performance under highly adverse conditions and prediction of system-level performance. Predictive models of material performance must include the physical impacts of temperature, pressure, and neutron fluence, as well as the complex chemistry occurring in water-based reactor cores. System-level reactor models incorporate the interactions among thermal hydraulics, neutronics, and structural mechanics.

As a specific illustration of the challenges faced in this research area, consider the challenges encountered in current power uprate and life extension programs for existing nuclear power plants. In the former case, scientists are interested in the performance of existing designs for fuel rod assemblies and reactor core structure to the increased physical stresses resulting from the uprate. These issues are largely related to materials performance under the combined insult of increased temperatures, pressures, and neutron fluence. Prediction of materials performance under these conditions requires understanding the coupled thermal hydraulics, neutronics, and materials response in dynamical simulations of the behavior of the core both at the local and the systems level. In the case of plant life extensions, scientists are further concerned with the prediction of materials performance on time scales beyond those that have been directly experimentally verified. This requires, for example, achieving sufficient basic knowledge of material performance subjected to neutron-initiated damage on 60- to 80-year time frames.

The goal for this research direction is a verified and validated science-based predictive capability for nuclear power plant construction and operation, including the full nuclear fuel cycle. This would include

the capability to successfully model and predict reactor performance under conditions of severe accidents (e.g., loss-of-coolant-type reactor incidents similar to the incident at the Three Mile Island Nuclear Generating Station in Pennsylvania), fuel performance under a broad range of operating conditions, and reactor performance under a power uprate and life extension program. This capability would allow the optimal refurbishment of existing plants and enable the design of entirely new systems. In all cases, acceptance by the U.S. Nuclear Regulatory Commission (NRC) of these verified and validated predictive modeling tools may reduce the current demand for extensive, costly experimentation.

Improving the safety, security, and economics of nuclear fission is a prerequisite for a revival of the nuclear energy market in the United States. Such a revival would affect national security on many fronts. First, it would lead to substantial cost savings resulting from better-defined safety operating and construction margins and would substantially decrease time scales for NRC construction and operating licensing. Second, it would reduce exposure to foreign manipulations of energy markets, especially in the area of fossil fuels. Third, it could have a substantial impact on the nation's ability to benefit from the present large stockpiles of fissile materials derived from both nuclear civilian and military sectors. These stockpiles could be viewed as national resources by providing the capability for either economical burning down-blended fissile material in light water reactors or in fast neutron spectrum facilities as part of a closed nuclear fuel cycle.

Carbon-Based Electricity Production

North America has access to abundant fossil fuel energy resources, principally consisting of recoverable coal, oils, and gases (the latter two mostly from tar sands and shale formations). Use of this natural resource is balanced by the need to sharply reduce the nation's greenhouse gas emissions, including that of CO_2. Because use of the aforementioned carbon-based energy sources does not recycle carbon, but rather represents a net CO_2 input to the atmosphere, it must be constrained by a combination of both increased efficiency in fuel use and in the sequestration of the emitted CO_2 (where feasible). Both increased efficiency and sequestration are technology driven and require a broad range of improved simulation models. Increasing efficiency requires advancements in basic material science and combustion chemistry. Developing a sufficient technical basis to consider geological sequestration demands simulation capability, including detailed material transport properties through potentially porous and/or fractured media. In all of these cases, predictive modeling capabilities in the areas of material science and the fluid dynamics related to combustion and flows through porous/fractured media will be key to success.

Advanced simulation capability has the potential to improve the efficiency of combustion systems while minimizing environmental impacts. Improving efficiency requires an understanding of compressible combustion in the context of high Reynolds number multiphase fluids. Modeling such flows requires high-mesh resolution for simulations to capture the relevant length scales present in the system. Optimization of system performance also requires details, including behavior of fuel sprays and material characteristics of surfaces exposed to high temperatures and pressures. Modeling sequestration requires developing a fundamental understanding of how rocks fracture, and how both liquids and gases are transported through media (e.g., rocks) that are potentially both porous and fractured. While the Reynolds numbers in this case are rather small, the computational complexity nevertheless remains high because the local transport coefficients are spatially highly variable.

Optimizing the use of fossil-based fuels for energy production and transportation needs would clearly minimize needs for such fuels at any given energy and transport demand level. Minimizing exposure to foreign manipulations of the global fossil fuel market significantly improves the nation's energy security.

Electrical Distribution Grid

The central problem of electricity distribution relates to the inherent mismatch between the production of electrical energy (dominated by large centralized power plants) and the use of that energy (dominated by highly time-dependent, distributed consumers). This is the "dispatching" problem faced by all electric utilities, a problem made considerably more complex by unpredictable supply disruptions resulting from power plant failures or transmission disruptions. The challenges faced in delivering improvements to distribution involve both economics (e.g., optimizing the balance between power produced and consumed, thus minimizing waste) and assurance and consistency of supply (e.g., suppressing grid instability). The technical challenges revolve around the capability for real-time monitoring and control of the grid. Improved understanding and modeling would enable evolution toward "smart grids" in which both energy providers and consumers are able to modify their behavior on the basis of grid energy measurements. Such a modeling capability would incorporate, for example, the time-dependent cost of electricity into account on both the supply and demand sides—a challenging synthesis of modeling both physics and economics. This requires that real-time measurement and assimilation of data streams from a large number of locations on the grid be coupled with models of overall grid response to provide real-time feedback to suppliers and consumers.

The key research task is to achieve a better understanding of grid stability and economics on the local and national scales using a hierarchy of increasingly complex (and increasingly highly coupled) grid models. Research in modeling this type of network interaction is made more challenging by a number of factors, each of which requires its own simulation model. The current electrical grid has substantial geographical differences in performance; e.g., models would be required to incorporate the age of various portions of the national grid. Sufficiently accurate models would allow design of upgrades on local, regional, and national scales that optimize distribution capabilities and minimize disruptions. As a practical concern, the ability to instrument the grid is highly nonuniform based on local conditions of available funding and interest. Ideally, a simulation model would provide input on optimal sensor deployment to maximize the national impact of local decisions. Ultimately, the grid simulation model must account for the nonlinear feedback as a result of the behavior of both energy suppliers and consumers as they react to the signals available from a fully instrumented distribution network.

Significant progress in predicting the performance of the national electricity distribution grid and in optimizing its performance are key steps in bringing about stable and secure power distribution. As the United States moves to an electricity-based transportation economy, based on a combination of plug-in hybrids and fully electric vehicles, reliability of electricity supply will become increasingly important to the nation's energy security.

Nuclear Fusion

In the long-term view of energy supplies, nuclear fusion will likely play a role in the latter part of the 21st century and beyond. While the successful practical deployment of nuclear fusion has the potential to provide an energy source that is sustainable over millennia, the challenges are formidable (Bromberg 1982). Consider the approach to nuclear fusion via plasma confinement by magnetic fields (inertial

confinement fusion raises drastically different, but no less difficult, challenges). Successful deployment of a magnetic fusion power reactor will require a sufficiently deep and predictive understanding of plasma stability to enable the design of active feedback for control systems. Similarly, predictive capability will be needed in material science to design a structure that can withstand the heat loads and neutron fluencies expected in fusion power. Each of these research areas presents multiphysics modeling challenges. A real-time control system requires a predictive simulation model with these and other physics models coupled.

Progress in fusion research has for some time relied on a combination of theory, computational modeling, and experimentation. Typically, progress has focused on relatively narrow problem areas (e.g., the performance of particular subsystems of a toroidal device) due to limitations of fundamental understanding and funding. It is now recognized that both theory and simulations must provide predictive capabilities at the systems level; e.g., assessment and real-time predictive capability of laser-plasma instabilities in inertial confinement fusions systems. Such systems-level modeling capabilities require predictive capabilities for the burning plasma couple with the various control subsystems of the fusion device. Several of these multiphysics application areas demand exascale computing (e.g., plasma instability, design of materials) to reach a predictive capability at a subsystem level. A practical engineering system would also incorporate computational fluid mechanics and control theory.

Success in this area would deliver a validated predictive capability for a viable magnetic fusion energy power reactor. Assuming the economics of such a power reactor are feasible, this would lead to the construction of the first magnetic fusion power reactor, thereby fundamentally changing the international energy market.

Success in making magnetic fusion energy reactors a practical reality will have enormous positive effects in protecting our nation's standing as a world-leading industrial power. The requisite fuel would be readily available within the territory of the United States. Progress in nuclear fusion is one of very few instances that would not only promote U.S. energy security, but also U.S. energy independence.

Critical Resource Protection

The United States increasingly depends upon arrays of remotely operated sensor and communication systems to gather and transmit information for a wide variety of critical issues, ranging from national defense and security to weather prediction. In the past few years, dependence on these advanced electronic systems is particularly evident in military intelligence and communication as well as Internet commerce. This dependence leads to increasing focus on the vulnerability of these networks to either local or system-wide failure. To protect these critical resources, a fundamental understanding of the threats is needed, along with the ability to model the underlying physical processes.

During the sequence of atmospheric nuclear tests conducted in the early 1960s, two new consequences of High Altitude Nuclear Explosions (HANE) became apparent: Electro-Magnetic Pulse (EMP) and Trapping of Energetic Particles (TEP) in the Earth's magnetic field. The physical processes in TEP include the phenomena often referred to as "belt pumping" in the literature. These consequences became evident for precisely the same reasons that still command attention: EMP disrupted electric power grids and TEP disrupted the satellites then in orbit (Glasstone and Dolan 1977; Northrop 1996; Foster et. al. 2004; Kennedy 2008; Schneider 2007; Gaffney 2005; McIntyre and Tussing 2008; Thomson 1995b).

The first evidence of these two effects was discovered during atmospheric tests conducted in the early 1960s. These effects were generally unexpected; hence, the underlying physics mechanisms were not diagnosed. In addition, tests were conducted with military secrecy in remote areas of the Pacific and the Union of Soviet Socialist Republics. As a consequence, there is a general lack of relevant data resulting in large uncertainly in the magnitude and fundamental mechanisms for the observed effects. The result is that military guidelines established to protect critical assets are exceedingly conservative and full compliance with these guidelines is prohibitively expensive. The guidelines are sometimes established as a compromise between desired levels of protection and what is fiscally possible. A more fundamental understanding of the multiphysics processes and the impact of these threats would enable prioritized protection criteria and anticipate emerging threats.

A pressing concern in the current political environment involves the asymmetry of these vulnerability issues (EMP 2004). Third-world advisories and terrorist organizations often depend on relatively little infrastructure of their own. Thus, while a HANE may indiscriminately affect every nation's satellites and electric power grids, these adversaries would be relatively unaffected.

A deeper understanding of these issues is needed. As with all nuclear testing, additional data from the full environment is not an option. Considering the energy densities and the geometric complexity involved, laboratory tests have limited applicability. Similar to the approach taken by the Stockpile Stewardship Program, laboratory experiments can be designed to validate portions of the relevant physical processes. Computational models are then used to predict performance in the full environment, extending complexity in geometry, scale, and multiphysics interaction.

The original numerical modeling of HANEs relied on fluid physics, but it was understood the particle mean-free-path was many times the scale of the nuclear fireball. In such cases, fluid models are inappropriate—giving answers that fail to account for the interaction of particle orbits and drifts well known to physicists who work with magnetized, collisionless particle kinetics. Although these limitations were well known, previously computational models containing the necessary physics fidelity were prohibitively expensive. With the introduction of terascale and petascale computing more appropriate, kinetic models are becoming feasible. For example, recent work has uncovered whole new paradigms of instabilities in the magnetic bubble that the expanding, conductive bomb debris stretches in the Earth's magnetic field (Brecht 2009). This new phenomenon is compatible with the limited observational data that exists but wholly unexpected from and nonexistent in the fluid models used previously (see Figure 2).

Exascale computing enables modeling of EMP and TEP with sufficient geometric fidelity and appropriate physics models to understand and protect against system-scale threats to critical resources. For both EMP and TEP, the details of the weapon's operation and the effect of these details on the weapon's photon and debris output are known to be important in the global-scale result. For example, the so-called E1 type of EMP involves gamma rays that generate free electrons that then gyrate in the Earth's magnetic field. Such simulations require resolution of the variations in the Earth's atmosphere and its magnetic field, measured in tens of kilometers, while modeling gamma ray productions requires submillimeter resolution. Similarly for TEP, the spectrum of the X-ray output and the charge state of the debris ions—a strong function of the details of the weapon disassembly—play a crucial role in the ultimate shape of the magnetic bubble. The shape of this bubble ultimately determines the degree of threat posed to satellites, even those located at higher altitudes. A comprehensive model must capture processes from the physics of the debris charge state within the bomb to the interaction of the electrically charged debris with the

global-scale, dipole field of the Earth. Exascale computing is required to bridge the tremendous gap in time and length scales that is required for the connection between detail and global effects.

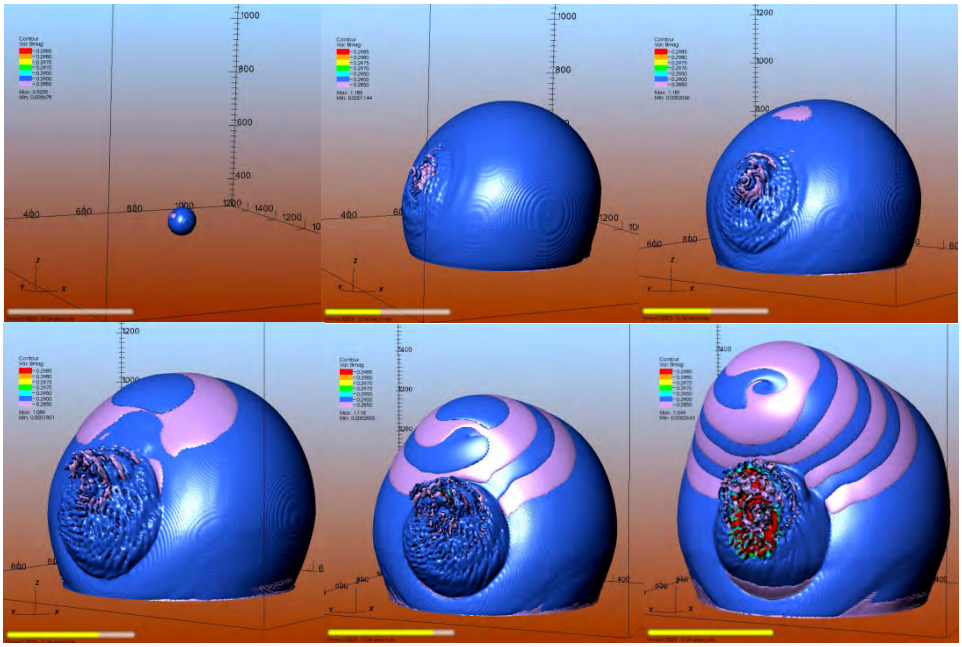

Figure 2. Evolution of a magnetized bubble in an inhomogeneous ionosphere due to a high altitude nuclear burst. The "helical mode" emanating from the bubble requires a computationally intensive kinetic ion model that cannot be represented using fluid model approximations. Image courtesy of D. Hewett (Lawrence Livermore National Laboratory).

Such an integrated simulation capability running on exascale computers would offer the insight that more complete simulations provide when compared directly to the limited observations of the past. This capability would also enable new ways of configuring laboratory high-energy density experiments to better understand unexplained observations. Finally, with experience comes increased confidence in the validity of these more complete models and the understanding gained as new phenomena (such as the helical mode shown in Figure 2) are discovered. The first step in preventing technological surprises is to know what is possible.

Enhanced Industrial Competitiveness

U.S. national and economic security depends on having innovative and agile manufacturing capabilities. In an industrialized world—widely described as "flat"—U.S. manufacturers must compete on technological differentiation, not cost, if the United States is to retain its position as one of the preeminent manufacturing countries.

HPC modeling and simulations are fundamental to the new way products are designed with virtual engineering and tests replacing costly, ineffective physical processes. Employing computational models enables rapid, far-reaching innovations not otherwise possible. Use of modeling and simulation also reduces development costs, certification costs, re-engineering costs, design cycle times, and improves

performance and efficiency while reducing waste such as emissions, noise, and raw material use (Council on Competitiveness 2008a, 2008b).

In the pursuit of their missions, DOE, the U.S. Department of Defense, and several other federal agencies have invested heavily in research and development for HPC technologies conducted by American HPC manufacturers, as well as academia and research laboratories. These investments have made the United States a world leader in HPC modeling and simulation, and in the design and manufacture of HPC systems. U.S. companies are responsible for producing a majority of the HPC systems appearing on the most recent list of Top 500 Supercomputer sites.

The use of HPC modeling and simulation has provided a competitive advantage for many of the manufacturing Fortune 50. However, officials from other countries realize the importance of simulation-based manufacturing and are working to "catch up." The drive toward exascale computing is vital if the United States is to retain its competitive edge and maintain its global leadership in manufacturing.

Application of High-Performance Computing Modeling and Simulation to Complex Industrial Systems

Modeling of complex industrial systems invariably involves statistical methods and engineering judgments whereby the governing equations are averaged. Direct numerical simulations of complex industrial systems such as high-performance gas turbine engines is simply not possible; however, in most practical cases, engineers need not simulate the flow through a gas turbine combustor to produce useful information, nor do they have the resources to analyze such detailed data even if it were available. The ad hoc models introduced during the averaging process necessarily result in a loss of information and confidence. To recapture what is lost, adjustable coefficients and terms embedded in these cryptic correlations and modeled equations must be supplemented with physical data derived from analysis, experimentation, or simulation. Simulations of complex industrial systems are therefore only as good as the models they contain (Moin and Kim 1997).

With the dramatic increase in the power of computing, scientists and engineers are now able to perform predictive simulations of an increasingly greater proportion of the fundamental physics from which all complex industrial-scale engineering models are derived. This capability enables the scientific community to reduce the range of scales that are crudely modeled by fine-tuning existing models and building improved closure approximations to resolve critical technical challenges that arise as the operability and performance thresholds of complex engineering systems are driven to their fundamental limits. In the past three decades, many industries have invested in leadership-class computing resources and embraced the role of predictive simulation science as a means to develop improved closure approximations for engineering models, as well as to supplement and inform costly experimental programs. Multipoint design, optimization and uncertainty quantification within an integrated computational design, analysis, and predictive capability framework are increasingly needed to resolve the many multiscale and multiphysics grand challenges that arise as industry strives to manufacture more efficient, reliable, sustainable and economic systems within increasingly competitive and regulated environments. Exascale computing for enhanced industrial competitiveness is a necessary component on the path to building such a predictive capability framework. The aerospace industry provides a key example.

Aerospace

Fundamentally, the grand challenge for the aerospace industry is to use extreme computing to design aerodynamically optimized aircraft including effects of turbulence, noise, propulsion systems, flexible structures, and active controls in a fully coupled manner (Peterson 1989). The importance for aerospace of modeling and simulation at the exascale is illustrated by examining three critical problems in the aerospace industry:

- airframe design
- noise mitigation
- jet turbine engine design.

As shown in Figure 3, virtual wind tunnels for airframe design that use high-fidelity simulation methods on today's fastest machines are restricted to single-point, wing-level assembly calculations. With the advent of exascale computing, full aircraft simulations at the required level of fidelity for predictive analysis with multipoint optimization across several thousand simulations will be possible. The transformational impact this advance in computing power will have is evident from the way HPC has fundamentally changed the way The Boeing Company (Boeing) has designed flight vehicles over the past three decades, a period that has experienced a comparable rise in available computing power to that projected on future exascale computing platforms (Garrett 2006).

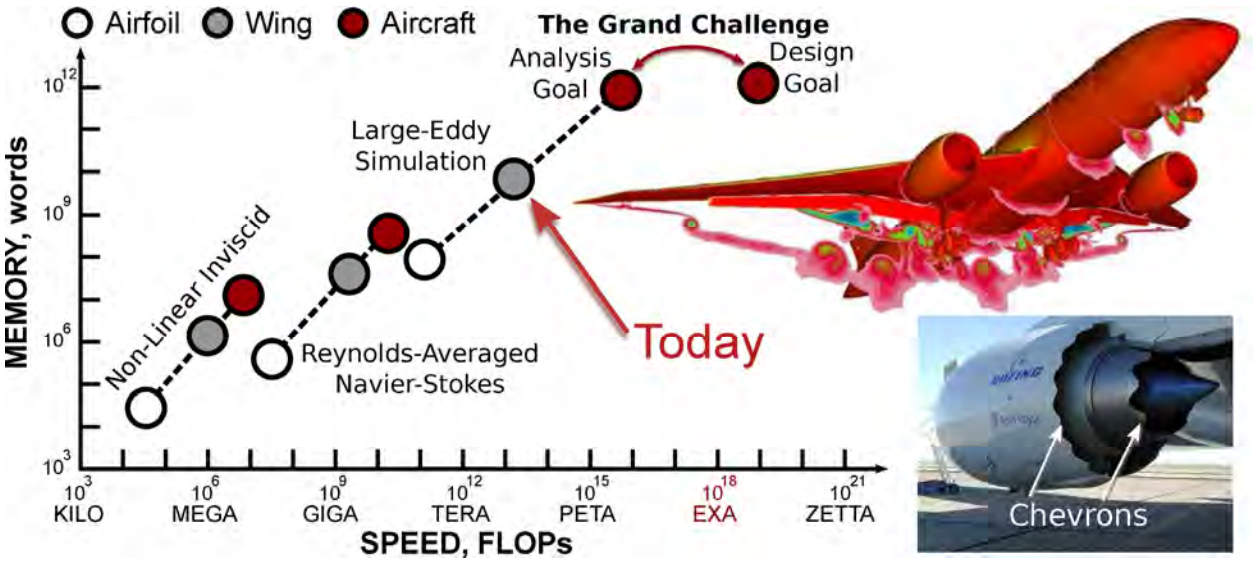

Figure 3. Left: Computer speed and memory requirements for the grand challenge of aerodynamic design. Right: Full Boeing aircraft simulation used to predict acoustic emissions during takeoff, highlighting the role of high-performance computers in providing solutions to the jet noise problem such as the design and optimization of chevrons. Image courtesy of Curtis Hamman, Stanford University. Sources: Garrett (2006), Chapman (1979), and Peterson (1989).

The evolution of HPC has enabled Boeing to create numerical simulation tools to assess system performance in nearly all of its products, including commercial aircraft, unmanned aerial vehicles, spacecraft, and military command platforms. This investment has led to the development of tools to

predict the aerodynamic performance of entire airplanes, optimize structural loads and minimize weight, and reduce the radar cross section of stealthy vehicles. Boeing has demonstrated that HPC-enabled design tools are good for business and provide for enhanced technology validation for application into new product development at lower overall cost.

With the widespread adoption of HPC into its design cycle, Boeing has been able to reduce the number of wing designs tested in wind tunnels from 77 in 1980 for the Boeing 767 airplane to only 11 in 2005 for the Boeing 787 airplane at a considerable cost savings. Rapid virtual prototyping, combined with virtual wind-tunnel simulations, has allowed Boeing to deliver improved performance and develop reduced-order models for airframe design optimization. In the future, Boeing envisions the availability of exascale computing will further enhance its design environment so that as few as two or three wing designs will require wind-tunnel testing, thus allowing Boeing to bring a new product to market in significantly less time and be more responsive to market demand. This exemplifies the dramatic reduction in cost and enhancement of industrial competitiveness made possible by extreme-scale computing.

Reducing aircraft noise is another great challenge confronting the aerospace industry (Garrett 2006). The design and development of both passive controls—such as Chevron Technology Ventures—and active flow controls to reduce aircraft engine noise has been driven by HPC, an example of which is shown in Figure 3. For example, Boeing has been able to simulate the noise reduction characteristics of multiple Chevron Technology configurations to select the best configuration for airport noise reduction before ever performing a wind-tunnel test or actual flight test. Even with today's computing capacity and simulation algorithms, there are still some flight conditions that require significant wind-tunnel testing to meet community noise requirements. This analysis is complicated by the fact that aeroacoustic noise generation covers a broad range of frequencies, multiple scales, and requires long-time integration to collect sufficient statistics of this very sensitive multiphysics problem. As extreme-scale computing becomes a reality, full aircraft aeroacoustic noise prediction, design, development, and optimization will be possible, which are key to improving the environmental performance and improving industrial competitiveness for the aerospace industry and related industries and national security applications where noise prediction and detection are of paramount importance.

In jet turbine engine design, computational fluid dynamics (CFD) was first used to complement component and engine testing in the eighties. Computer simulations provided valuable, detailed flow field information not typically available in experimental testing. As a diagnostic and troubleshooting tool, CFD found widespread early adoption in the combustor design community. As HPC systems became more readily available, the role of CFD in combustor design has fundamentally changed. Consider the impact HPC has had on the combustor design process at Pratt & Whitney (Figure 4) (Kim and Syed 2004; Reynolds et al. 2003). A large component of combustor development cost has historically been due to high-pressure rig tests where component tests are performed using actual engine hardware at realistic, elevated pressure levels. By the early nineties, the required number of high-pressure rig tests was reduced by 25% from that needed just 5 years earlier. This significant savings in cost and time was primarily due to the adaptation of reduced-order engineering models through early CFD of model problems and the diagnostic ability of CFD to supplement and probe the experimental tests in more detail.

As exascale computing capabilities become available, breakthroughs in operability, durability, and emissions of modern combustors will come through integrated systems-level design optimization as the above examples indicate. The combustion industry will therefore need tools able to resolve the higher-

fidelity details of turbulent mixing, chemical reaction, multiphase spray dynamics, and species transport present in fully coupled engine operation. This improved physics understanding, combined with refined geometrical and physical fidelity enables the development and adaptation of reduced-order models for parametric studies of combustor systems. From this, a streamlined optimization methodology for modern combustor systems will be possible, guided by the need to appropriately apply physics-based models of the required level of fidelity and thereby accurately predict and quantify the uncertainty in these virtual experiments. The transformational impact such a design tool will have upon the combustor design community is a direct result of the investment in extreme-scale computing.

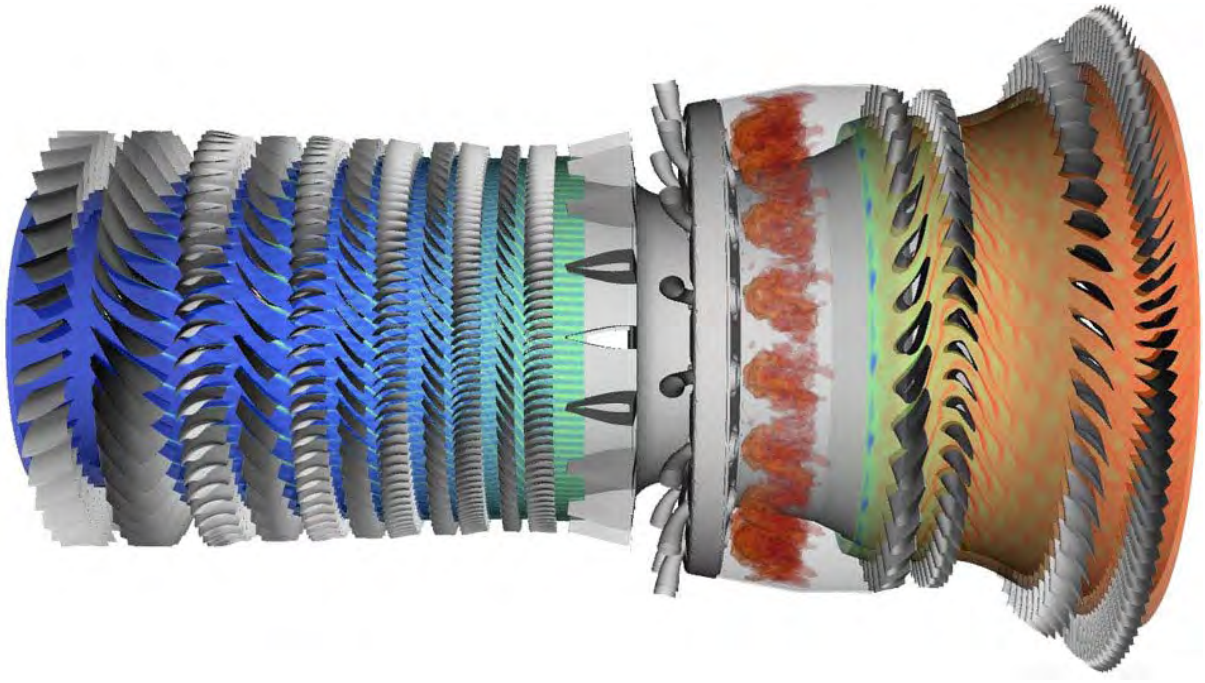

Figure 4. Comprehensive simulation of the flow through an entire jet engine performed by Stanford University and Pratt & Whitney. The compressor (left) and the turbine (right) are computed using unsteady Reynolds-averaged Navier-Stokes (URANS); the combustor (center) is computed using Large Eddy Simulation. Data are transferred across these interfaces using a coupled, multicode, multiphysics integration environment. Contours of entropy in the high-pressure compressor and in the first two stages of the turbine are shown while the flame in the combustor is visualized with an isosurface of temperature in a realistic Pratt & Whitney engine. Image courtesy of Center for Turbulence Research, Stanford University.

As complex engineering systems are driven to their fundamental limits, more accurate computational analysis tools to reduce the cost and schedule during the design, operation, and maintenance of these systems will be greatly needed by industry. Adopting physics-based design tools, combined with reduced-order models for engineering design optimization, are urgently needed as soon as such tools and models are practical. The development and deployment of extreme-scale computing architectures will herald the coming of a fully coupled, system-level optimization and engineering design cycle that will greatly enhance industrial competitiveness and national security science.

CONCLUSIONS

Exascale computing will offer significant advances in many fundamental areas, including chemistry, material science, and nuclear physics. Improved fundamental understanding of these and other topics is the foundation on which advances in multiphysics simulations are built. The scientific community anticipates insight into the behavior of nonlinear coupled systems with an unprecedented ability to predict key observables. Progress will enhance traditional NNSA application areas, such as the successful Stockpile Stewardship Program, and enable the capability to expand its national security mission into protection of critical infrastructure. Progress will also enhance traditional DOE Office of Science application areas, such as energy security. Through collaborations with industrial partners, these capabilities can increase national competitiveness in industrial engineering and manufacturing arenas.

One of the grand challenges facing scientific computation is the understanding and propagation of uncertainties in simulations. These uncertainties may arise from fundamental data input into simulations, limitations in physics models, or inadequacy in understanding of the physical processes. These challenges are particularly complex in multiphysics simulations, where scientists must distinguish between inherent uncertainties and physical processes omitted from the simulation. They are amplified yet further in cases—such as certification of the nuclear stockpile—in which systems-level experimental validation is either difficult or impossible to obtain. As multiphysics simulation results reach fidelity levels that make them valuable tools for decision-makers in industry and government, advances in quantification of uncertainties will become drivers for exascale computing.

The recent vast increase in available computing power resulting from innovations in massively parallel computer architectures makes it possible to obtain accurate solutions for some problems of interest with extant algorithms. It remains the case that full-fidelity calculations for the most challenging multiphysics problems will require further algorithmic innovations, as well as concerted efforts to design new types of validation experiments and match results with the highest fidelity simulations.

Multiphysics challenges are present in nearly every aspect of the DOE mission. Recent advances in computing, theory, and experimental techniques have led to new predictive capabilities for some noteworthy multiphysics challenges. Exascale computing will enable predictive capability with quantified uncertainties for increasingly complex multiphysics conditions extending to the systems in which these conditions manifest. This computing capability is required to deliver the needed fundamental understanding within the DOE mission and to enhance national and global security.

NUCLEAR PHYSICS

Co-Leads: Joseph Carlson, Los Alamos National Laboratory
Edward P. Hartouni, Lawrence Livermore National Laboratory

Panel Members: Aurel Bulgac, University of Washington; Wick Haxton, University of California, Berkeley/University of Washington; Calvin Johnson, San Diego State University; Robert Little, Los Alamos National Laboratory; Tom Luu, Lawrence Livermore National Laboratory; Petr Navratil, Lawrence Livermore National Laboratory; Ian Thompson, Lawrence Livermore National Laboratory; James Vary, Iowa State University; Robert Wiringa, Argonne National Laboratory; Walid Younes, Lawrence Livermore National Laboratory; Witek Nazarewicz, University of Tennessee; Esmond Ng, Lawrence Berkeley National Laboratory; Erich Ormand, Lawrence Livermore National Laboratory; Sofia Quaglioni, Lawrence Livermore National Laboratory; and Kenneth Roche, Pacific Northwest National Laboratory

CURRENT STATUS

Nuclear science and the energy produced by nuclear processes are at the core of the National Nuclear Security Administration (NNSA) program and central to its mission. Nuclear reactions are critical in many applications, including the National Ignition Facility (NIF), energy production, weapons, and global threat reduction. Nuclear reactions are the source of energy in all of these applications, and they can be crucial in understanding and diagnosing the complex energetic environments integral to NNSA's work.

Nuclear processes are complex quantum many-body problems. Modeling and simulation of nuclear reactions and their role in applications, coupled with experiments, have played an integral role in achieving NNSA's mission. The science input to NNSA program applications has been heavily reliant on experiment combined with extrapolations and physical models "just good enough" to provide a starting point to extensive engineering that generated a body of empirical information. The NNSA program requires reliable extrapolation of these models into domains beyond the existing data sets. Additional requirements exist to quantify the uncertainties within the regime of existing data. These requirements arise at a time when the ability to perform additional experimental tests is no longer possible, especially tests that produce data in the extreme environments that uniquely characterize these applications. The end of testing has required improvements of the predictive capabilities of codes simulating the reactions and associated applications for both well-known and well-characterized cases, as well as incompletely known cases.

Developments in high-performance computing (HPC), computational physics, applied mathematics, and nuclear theory have combined to make spectacular advances in the theory of fission, fusion, and nuclear reactions. Current research exploits these developments in a number of U.S. Department of Energy (DOE) Office of Science and NNSA programs and in joint programs such as Science Discovery through Advanced Computing (SciDAC) that supports the project, *Building a Universal Nuclear Energy Density Functional*, whose goals are to provide the unified approach to calculating the properties of nuclei. The successful outcome of this and similar projects is a first step toward a predictive nuclear theory based on fundamental interactions between constituent nucleons.

The application of this theory to the domain of nuclei important for national security missions will require computational resources at the extreme scale, well beyond what will be available in the near future.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

Recent developments in theory and computation make it possible to project that calculations will be made that describe the details of nuclear reactions relevant to NNSA's program. The calculations will be performed with the highest physical realism starting from the correct fundamental description of the nuclear physics, providing true predictive capability with regards to nuclear phenomena. Exascale computing will make these claims realizable.

Based on this recent progress, the panel identified four major priority research directions (PRDs) that will be achieved in nuclear science using extreme-scale computing over the next decade. The first PRD is the **physics of fusion**, which includes computing the properties of nuclei, their interactions, and cross sections of essential reactions that both generate energy and enable crucial diagnostics in fusion environments such as those found in inertial confined fusion (ICF) (e.g., the NIF) and nuclear weapons and astrophysical phenomena. The second PRD is the **microscopic description of fission**, which includes computing pre- and post-scission quantities and providing a consistent picture of the many types of fission data (cross sections, fission product properties, prompt fission neutrons, gamma rays, etc.). These processes are critical for energy; global security; and weapons applications, including stockpile stewardship. The third PRD is the **theory of neutron-induced reactions**, which will provide an accurate, predictive capability from first-principles calculations. The fourth PRD is the **simulation of integrated nuclear systems** to provide necessary physics realism to other applications—such as nuclear reactors—that require the same types of nuclear physics input as the NNSA applications. High-fidelity simulations of these problems, including accurate particle transport, reactions with "complete" correlations and uncertainties, and nuclear burn, require dramatic improvements of current computational capabilities. Simulations of these integrated systems, either fission or fusion reactors, will be a necessary component of exploring possible future energy sources.

PRIORITY RESEARCH DIRECTIONS

Physics of Fusion

Fusion reactions in light nuclei are critical for both basic science and NNSA applications. These reactions aid in understanding the early universe and provide the energy that powers the stars. These reactions are also essential elements of NNSA science. To diagnose the performance of nuclear weapons and fusion at the NIF, a clear and precise understanding of these reactions is critical. While the basic features of fusion have been identified for some time, crucial issues must be resolved to take full advantage of experimental results and extend understanding to regimes that cannot easily be measured.

At present, R-matrix analyses of experimental data are the primary source of information used in multiphysics codes for NNSA. These analyses are closely associated with nuclear experiments and require a rather extensive data set to fully characterize the reactions. The goal of large-scale computations over the next decade is to reliably calculate these reactions starting from realistic nuclear interactions. Validating these calculations against well-known data will lead to a predictive capability

that will allow the prediction of reaction rates that are difficult or impossible to measure because of unstable targets or extremely low reaction energies and to predict rates in hot and dense environments.

Several methods have recently been developed to accurately calculate the properties of light nuclei from a realistic nucleon-nucleon interaction. These algorithms—Green's function Monte Carlo (GFMC), the No-Core Shell Model (NCSM) coupled with the Resonating Group Method (RGM) for reactions, and coupled-cluster (CC) methods—take advantage of the best available computational resources and have enabled developers to make great strides over the past few years (Pieper and Wiringa 2001; Navrátil et al. 2000, 2007; Hagen et al. 2007b). These methods have provided an accurate, microscopic understanding of the structure of all nuclei up to carbon and closed-shell nuclei, including calcium. This is a significant achievement because the structure of light nuclei is exceedingly complex, with significant many-body correlations and cluster structures. For example, one of the lowest excited states of carbon (the "Hoyle" state) is nearly degenerate with the energy of three isolated alpha particles. This property is necessary to allow just the right amount of carbon to be created in stellar environments.

To be useful to NNSA applications, realistic *ab initio* approaches must tie these calculations of nuclear structure to the calculation of low-energy reactions (Figure 5). These calculations are still in their initial stages, but dramatic successes are emerging (Nollett et al. 2007; Quaglioni and Navrátil 2008; Hagen et al. 2007a). For example, calculations of n-alpha scattering now successfully match all of the low-energy s- and p-wave phase shifts, including the positions and widths of the low-lying resonances.

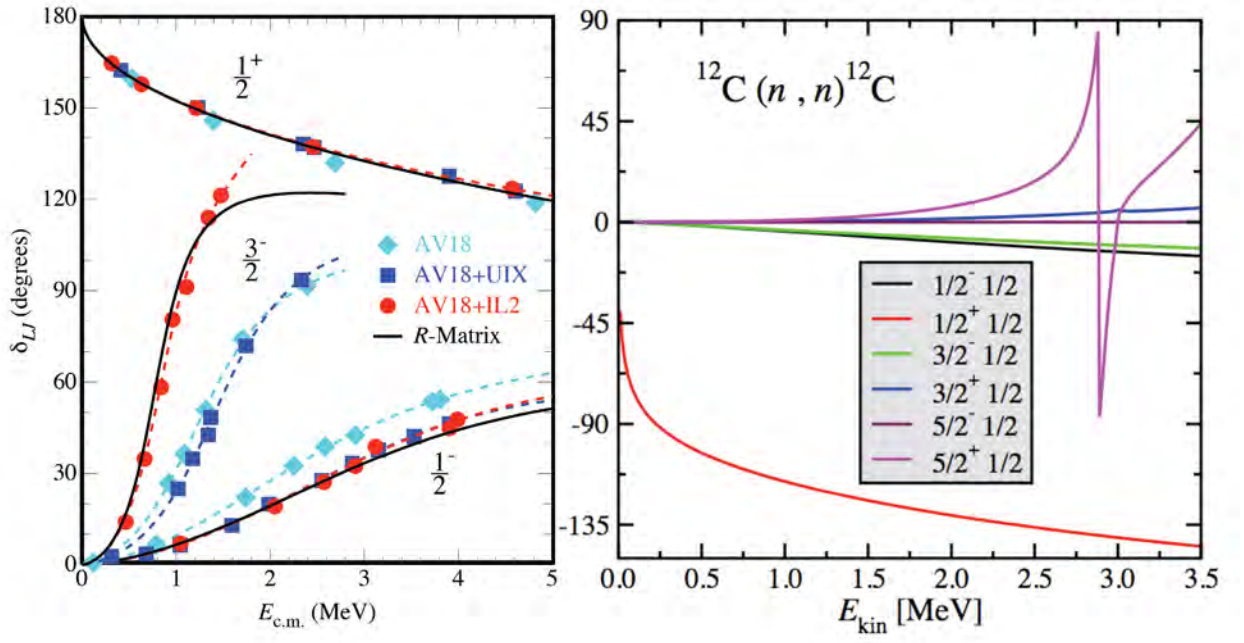

Figure 5. *Ab initio* calculations of low-energy nuclear scattering, low-partial wave phases versus energy. The left panel compares neutron-alpha scattering calculations with the R-matrix analysis of experimental data. Recent theoretical progress and computational advances are making it possible to calculate neutron scattering on complex nuclei such as carbon (right panel). Image courtesy of Joseph Carlson (Los Alamos National Laboratory) and Edward P. Hartouni (Lawrence Livermore National Laboratory).

Researchers are on the verge of calculating thermonuclear reactions in *ab initio* approaches, but they are severely limited by presently available computing power. Currently, the *ab initio* calculations of nuclear properties tax the largest available computers. Recent calculations of the carbon ground state required several million core-hours, even taxing the available resources running efficiently on up to several hundred thousand cores. Calculations of even simple reactions will be far more complex because they involve properties of two or more nuclei and their mutual interactions. They will require one to three orders of magnitude more computing power than structure calculations in the same mass region.

Initial calculations will be for d-t reactions and other cases where precise experimental data already exist. This will allow for refinement of the algorithms and more precise tuning of the three-nucleon interaction than has been possible to date. A suite of such calculations will require an order-of-magnitude increase beyond current capabilities and can be accomplished in 2 to 3 years. In addition to computational issues, sufficient labor resources will be required to pursue this work. While initial algorithms have been identified and implemented, much remains to be done to enable these codes to work efficiently and accurately over a range of energies.

The scaling of GFMC calculations with nucleus size is well understood and can reliably estimate the increase in computer resources needed to address the ground state of the ^{16}O nucleus. Presently, the GFMC calculation of the ^{12}C ground state requires approximately 400 peta operations. The number of operations will increase by a factor of approximately 1200 for ^{16}O, with growth provided by the available computing resources increasing from the petascale to the extreme scale. Further increases of one to two orders of magnitude are necessary to go from ground-state structure calculations to scattering and reactions.

The present ability in GFMC was obtained by splitting the work on one Monte Carlo configuration (one set of coordinates, including all spin and isospin states) among tens of cores. Until recently, only one core was used. For mass A=16, the work will have to be shared at an even finer level. Many cores will have to work on the computation of one wave function, and, because of memory limitations, operations involving wave functions stored on different nodes will be necessary. Because of the branching random walk nature of the GFMC algorithm, future advances also will require the efficient use of many, more loosely coupled nodes, which is being addressed with the introduction of the Asynchronous Dynamic Load Balancing Library (ADLB).

In addition, the scaling of NCSM/RGM scattering calculations with nucleon number can be estimated. Currently within the NCSM/RGM framework, it is becoming feasible to calculate low-energy nucleon-^{12}C or nucleon-^{16}O scattering that take into account a few lowest excited states of the target nuclei with soft two-nucleon (NN) forces using approximately 1000 cores on present-day machines. Deuteron–alpha and deuteron-^{3}H scattering calculations with target nuclei restricted to their ground states are of a similar scale. With the increase of the projectile energy, more channels must be included from excited states of the target nuclei. Furthermore, it becomes obvious that three-cluster channels are important for light-ion fusion reactions; in particular, those that involve the deuteron or triton—nuclei that easily break up. Proper treatment of three-cluster channels is a challenging, presently unsolved problem. Within the NCSM/RGM, the three-cluster channels can be included using the hyperspherical-harmonics basis expansion of the three-cluster final states and/or either binary or three-cluster state for the case of closed channels by coupling of the NCSM/RGM binary-cluster basis with the NCSM *A*-nucleon basis (for a reaction involving the total of *A* nucleons). The computational demand increases dramatically (a factor of approximately 10^6) with the increasing size of the projectile (from a single nucleon to an alpha particle),

with the increasing number of channels from excited states of the target (from 2 to 10), from the treatment of the three-cluster states, and by including the three-nucleon (NNN) interaction. Therefore, this clearly is a problem requiring extreme-scale computation.

The ground states of ^{16}O and ^{40}Ca can presently be computed within the CC method. Here, the inclusion of NNN forces is challenging, and estimates put its computational expense at the petascale. The computation of excited states is an order of magnitude more computationally expensive because of the proximity of the scattering continuum. It will be determined on a Gamow basis, consisting of bound, resonant, and scattering states. The lowest-lying excited 0+ state in ^{16}O is an alpha-particle excitation and requires the inclusion of four-particle, four-hole cluster configurations. The computational resources required for the calculation of this state are estimated to be at a scale of tens to hundreds of peta operations and can be performed on current and next-generation machines (up to 20-petaflop machines). Because of the growth of the number of cores by a factor of approximately 1000, it will be challenging to use an extreme-scale computer for these calculations. Figure 6 illustrates computational advances in fusion progressing toward exascale computing.

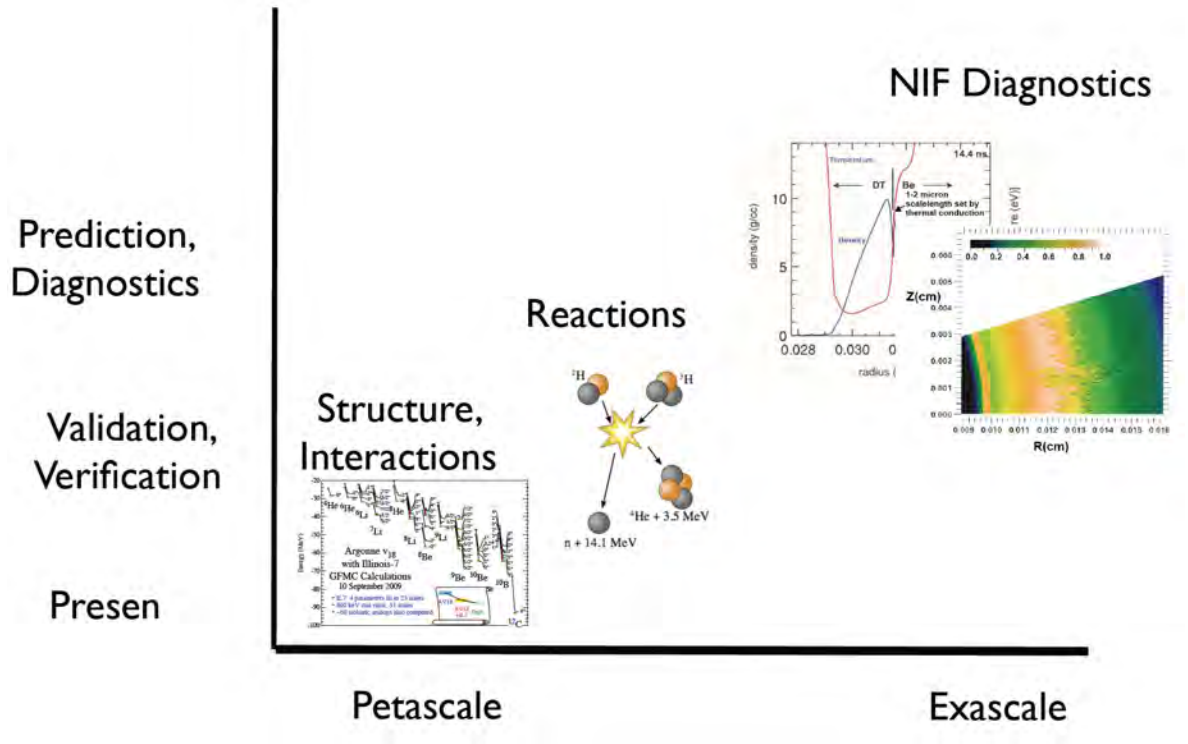

Figure 6. Computational advances in fusion progressing toward exascale computing. Present calculations are dominantly of nuclear structure, with progress towards scattering and reactions. Exascale computing could enable calculations that will help diagnose complex environments such as the National Ignition Facility. Image courtesy of Joseph Carlson (Los Alamos National Laboratory) and Edward P. Hartouni (Lawrence Livermore National Laboratory).

An essential input to all of the algorithms already mentioned is nuclear force. Nuclear force is itself a manifestation of quantum chromodynamics (QCD). Attempts to formally derive this force have been met with limited success. As a result, the forces that bind nuclei have been determined experimentally. In

turn, this result has limited the knowledge of certain key aspects of nuclear force that are not easily accessible experimentally, such as pieces of the NNN force. With extreme-scale computing, it will become possible to extract nuclear forces using computational techniques from lattice quantum chromodynamics (LQCD). When numerically calculated from QCD and used as input into the algorithms mentioned, these forces will allow one to connect the forces between quarks and gluons (QCD) to the forces between nucleons. This is an extremely challenging problem because of the vast range of scales involved, from the order of a gigaelectronvolt (GeV) scale of nucleon and nuclear masses, to the megaelectronvolt (MeV) scale for binding energies, to the less than one millielectronvolt (meV) required for nuclear reaction rates. In addition, the computational resources needed to simulate light nuclei via LQCD increase dramatically with the size of the nucleus (Luu and Walker-Loud 2009). Recently, progress has been made in studying NN systems on the lattice (Beane et al. 2006, 2008), providing a direct link between the QCD and the NN interaction. Already, there are initial studies of three-baryon systems, such as the triton. These calculations are the initial steps in calculating the NNN interaction, a quantity of great value to NNSA because of its effects on nuclear structure and reactions. More challenging still will be the calculation of reaction rates directly from QCD.

In any of these approaches, predicting important rates not yet measured will be extremely challenging. The most difficult problems—those with weakly bound initial or final states and multiple clusters—will require on order of 1000 times the largest available computing resources. These will allow treatment of many important problems that have not been addressed adequately in any framework, including effects of multifragment final states and reaction rates in hot and dense environments. Accurate predictions for such processes, particularly applications to light element reactions and synthesis, would be extremely valuable for NNSA, NIF, and more broadly for nuclear science.

The most important fusion reactions involve weakly bound particles in the initial state and tightly bound ones in the final state because that is how energy production is maximized. The weakly bound particles in the initial state present an especially difficult challenge in *ab initio* calculations because their spatial extent requires a much bigger range of scales to evolve the wave function to its known asymptotic form. In this regard, the quintessential reaction is d+t fusion in the ^{5}He system, but other reactions involving deuterons are also important. The experimental data are remarkably complete and consistent in the A=5 systems (^{5}He and ^{5}Li), making them good benchmark calculations to compare with the well-determined scattering amplitudes from R-matrix fits and illuminating tests of what are primarily Coulomb differences between the systems. Single-channel calculations of N+α scattering already have been quite successful in reproducing the R-matrix phase shifts, but extending those calculations above the d+3N threshold will require significantly more computing power.

Radiative capture reactions in these systems, such as $T(d,\gamma)^{4}He$ and $^{3}He(d,\gamma)^{4}He$, can also be used as diagnostics of fuel temperature in both inertial and magnetic confinement fusion. A detailed knowledge of the photon spectra in these reactions could be greatly enhanced by *ab initio* calculations, which is essential because experiments are hampered by the high background of energetic nucleons produced by the dominating particle-production reactions.

Other light systems of initial interest are ^{7}Li and ^{8}Be. Both systems, like those for A=5, have only a single channel open over a wide range of energies. However, above that range, they contain reactions with neutrons and deuterons on ^{6}Li that are of great interest in NNSA applications. Several fusion reactions relevant to the NNSA programs whose understanding will be greatly helped by *ab initio* studies have multibody final states. Those include $T(t,2n)^{4}He$ (the fusion of two tritium nuclei into an alpha

particle and two neutrons), which is important in any d+t fusion system (such as the NIF), and the reactions in the ^{9}Be system, ^{6}Li(t,n)2 ^{4}He and ^{7}Li(d,n)2 ^{4}He, producing a neutron and two alpha particles.

Multicluster (more than two) fragments are difficult to incorporate into the R-matrix or *ab initio* methods. The complete scattering matrix in such cases is quite involved and may be difficult to calibrate using an *ab initio* approach. However, approaches can certainly be implemented that respect specific scattering boundary conditions that should severely restrict the available solution space for the physical scattering matrix. Exploratory studies are now underway, and significant work is required to examine the optimum strategy for gaining the most information from large-scale simulations.

Exascale computing will play an important role in *ab initio* studies of T(t,2n)^{4}He; n-^{6}Li and n-^{7}Li reactions—in particular for higher neutron energies (14 MeV) and deutron-^{6}Li and triton-^{6}Li reactions—as well as seemingly simple processes such as neutron-triton scattering and breakup. The d-t fusion will also trigger γ rays from T(d,γ)^{5}He radiative captures. The high-energy photons produced by this process could be used to determine the ion temperature of the fuel in inertial or magnetic confinement fusion. A robust knowledge of these photon spectra is essential and can be greatly enhanced by *ab initio* calculations because experiments are hampered by high background of the 14 MeV neutrons produced by the dominating T(d,n)^{4}He reaction. Note that even in situations when a complete and exact *ab initio* description would not be achieved, *ab initio* calculations can provide guidance and input to the R-matrix analysis of experimental data.

It is also critical to understand how to bridge from the nuclear to the atomic scale to address the impact of hot and dense environments on nuclear reaction rates. Particularly in the case of very weakly bound systems, electronic screening effects should be investigated explicitly in *ab initio* approaches. Here, computational approaches can be exploited to study systems that are especially difficult to address experimentally. The huge range of distance scales involved would require an additional order of magnitude of computer resources. Initial studies on lighter systems, most importantly d-t, would measure the importance of these effects and the necessity of larger-scale calculations in heavier systems.

In summary, a predictive capability for fusion reactions is critical to advance the understanding of nuclear processes relevant to the NNSA and, in particular, to exploit the advances made in NNSA experimental facilities, including NIF. Prospects for *ab initio* approaches to nuclear reactions are extremely promising, and the combination of computational and theoretical advances will enable a dramatic impact for theory and simulations in this arena.

Nuclear Fission

Understanding nuclear fission at a more comprehensive level is a critical problem in national security with important implications in nuclear materials detection and nuclear energy, as well as enhancing the quantitative understanding of nuclear weapons. Fission also plays an important role in the r-process, the synthesis of all the heavy elements in the universe. Present-day fission models used to develop nuclear data libraries for all of these applications are quite simple, typically involving a few parameters that can be adjusted to reproduce the basic features of experimental data. Although use of fission models is a productive approach, it is time to implement more sophisticated models based on microscopic nuclear input. In short, in radioactive waste management, nuclear engineering and reactor physics, and stockpile stewardship, a high-quality nuclear-theory input is needed with quantifiable error bars.

A truly predictive model of all fission observables is an extremely challenging task. New approaches should be applicable over a wide range of energies and nuclear masses, as well as be able to make predictions about more detailed questions, such as the time and energy correlations of emitted neutrons. The SciDAC-2 Universal Nuclear Energy Density Functional (UNEDF) project (Bertsch et al. 2007) and petascale computing resources are opening the way for a comprehensive microscopic description of both the static properties of atomic nuclei and the fission process.

Nuclear fission is an N-body, quantum mechanical, large-amplitude collective motion (LACM) problem involving a large-scale collective rearrangement of nuclear matter. Microscopic approaches, such as the nuclear density functional theory (DFT), require very large computing resources. Recently, these approaches have been implemented on the most powerful available computers, running efficiently on machines with tens of thousands of cores. When combined with the physical insight gained through comparison of simulations with data, this advanced computing power has yielded significant progress in the theory of nuclear fission. In particular, static properties, such as spontaneous fission barriers and half-lives, can be described reasonably well (Staszczak et al. 2009), and significant progress has been made in compound-nucleus fission barriers (Pei et al. 2009; Sheikh et al. 2009), as well as fission-fragment properties for neutron-induced fission (Younes and Gogny 2009a, 2009b). The limitations typical of previous calculations (e.g., the presence of mean-field symmetries artificially constraining the problem) that crippled this type of calculation are starting to be lifted with advanced, present-day capabilities.

As a first step, adiabatic approaches to fission require the determination of the potential energy surface (PES) in a multidimensional space of collective coordinates, which comes from constrained Hartree-Fock-Bogoliubov (HFB) calculations (see Figure 7). A promising starting point is the nuclear DFT, which can explain and predict nuclear structure across the complete table of nuclides. The most accurate functionals currently in use are purely phenomenological and have parameters fit to a selected group of nuclear properties. Petascale computing and improvements in DFT codes are allowing improved calculations of the comprehensive set of nuclear properties for all nuclei. Potential energy surfaces for fissioning nuclei have been studied in macroscopic-microscopic models, which provide valuable guidance to microscopic approaches. Potential energy surfaces in microscopic approaches come from constrained HFB calculations (Staszczak et al. 2009). Including all relevant degrees of freedom to obtain a realistic and precise PES will be a particularly challenging task. In addition, the description of dynamics will require the evaluation of the multidimensional inertia tensor, both in the particle-hole and particle-particle channels (Baran et al. 2007).

A first step beyond current capabilities is to fully follow the adiabatic time-dependent (ATD) HFB theory that constitutes the limit of the imaginary-time approach (Skalski 2007). Spontaneous breaking of intrinsic symmetries requires central processing unit (CPU)-intensive calculations. At least four degrees of freedom—elongation, mass asymmetry, necking, and triaxiality—must be considered. In this collective space, the optimum collective that minimizes the collective action must be located to evaluate the barrier penetration probability, or a fission half-life, by integrating the action along this optimum path. On its path to fission, the nucleus passes through numerous saddle points in the collective space. In fact, saddle points obtained in calculations constrained by only one collective variable are sometimes incorrect. Hence, special numerical techniques are required to find them. To bring other constants-of-motion, such as angular momentum, to the problem poses significant challenges that will require new conceptual and algorithmic developments.

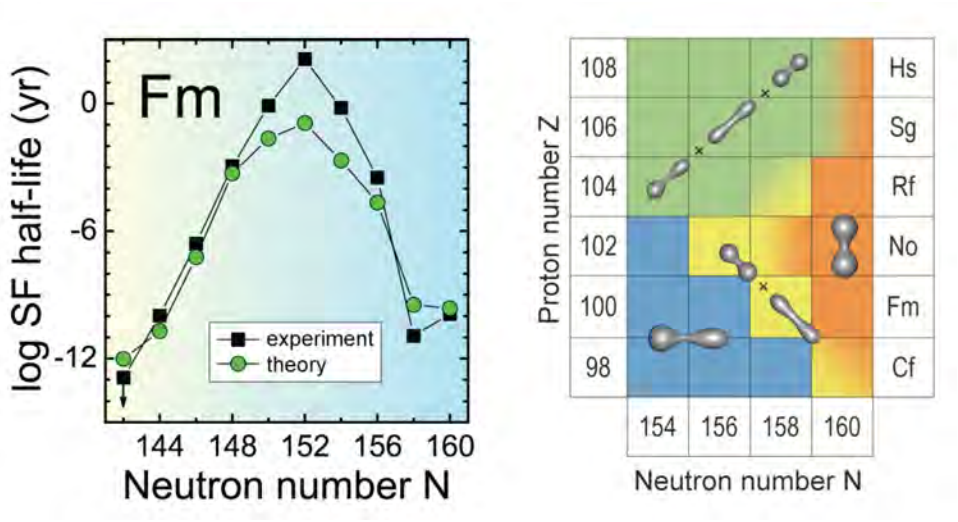

Figure 7. Fission calculations at present. The left panel compares experimental and calculated half-lives for fermium isotopes versus neutron number. The right panel shows the different fission products that arise as a function of neutron (N) and proton (Z) number. Image courtesy of Joseph Carlson (Los Alamos National Laboratory) and Edward P. Hartouni (Lawrence Livermore National Laboratory).

To attain sufficient mesh refinement, it will be necessary to compute on the order of 100,000-plus constrained HFB calculations for every single nucleus. Because a single HFB run with all self-consistent symmetries broken takes between 6 and 10 hours of CPU time, depending on how sophisticated calculations are, it takes 20 to 30 CPU-years to carry out the full fission pathway analysis for 20 isotopes. Hence, massively parallel computer platforms are required.

Time-dependent versions of such methods are needed to study the dynamics of the fission process and require even more processing power. Several approaches can be applied to the description of nuclear fission and will be pursued. Each entails a number of serious computational challenges.

The time-dependent generator coordinate method (TDGCM) is a fully quantum-mechanical formalism within which the dynamical and collective aspects of fission can be studied. In this approach, a wave packet is built from the microscopic HFB solutions obtained for a broad range of values of the relevant collective degrees of freedom. Then, according to the laws of quantum mechanics, the wave packet is allowed to evolve in time toward scission. The TDGCM already has been used on a limited computational scale to predict realistic fission times (Berger et al. 1984) and fission yields (Goutte et al. 2005) for thermal fission. These previous applications were limited to two collective degrees of freedom. The future of this promising technique lies in its extension to many more collective degrees of freedom (at least the four previously mentioned—elongation, mass asymmetry, necking, and triaxiality—along with fragment separation) and the coupling of collective and intrinsic excitations of the nucleus. The ability to couple intrinsic and collective motions of the nucleus will enable the description of higher-energy-induced fission (up to several MeV in incident neutron energy) by removing the adiabatic approximation inherent to current microscopic fission theories. With existing peta floating point operations per second (PFLOP/s) machines, calculations with this extended TDGCM would require approximately 1 year of continuous computations for one nucleus. The same calculations could be completed in less than a day for each nucleus. In its full exascale application, the TDGCM will produce

realistic predictions of fission yields, fragment properties, and fission emission spectra for important nuclei that cannot be studied in the laboratory.

The reactions of induced fission with neutrons and energetic photons are of central importance to the NNSA mission, particularly for nuclear forensics. To produce reliable theoretical predictions, a qualitatively new approach to nuclear reactions will have to be implemented—an approach based on a real-time, path-integral formulation of the quantum many-body problem (Negele and Orland 1987). Because of the extraordinary numerical complexity of this formulation, this approach has never been applied to a real physical situation, except some relatively simple models where the proof of principle has been produced. Relatively simple cases also have been studied in a related framework in the case of chemical reactions involving a small number of electrons and nuclei (Tully, Casida, Tavernelli, Sugino, and many others).

In the case of induced fission, unlike the very low-energy nuclear reactions that can be modeled with collective motion and adiabatic approaches, the role of quantum and statistical fluctuations dominates, changing the nuclear dynamics in a qualitative manner. The quantum fluctuations, in particular, turn quantum dynamics basically into a stochastic process; then, the nuclear system, instead of following a specific PES (as in the case of low-energy nuclear reactions), hops from one potential surface to another in a random manner governed by subtle quantum processes. Unlike the case of chemical reactions studied thus far, a nuclear system performs an extremely large number of quantum jumps. The large number of quantum jumps is responsible for the large number of final channels in an induced fission reaction. Until now, these processes have been modeled within various theoretical models, the validity and accuracy of which are essentially impossible to quantify.

However, the same large number of quantum jumps also leads to a marked irreversible character of the nuclear dynamics, which until now has been described only within primitive phenomenological models that neglect completely the quantum nature of these processes. With the advent of exascale computers, scientists can envision a qualitative new approach to the quantum many-body dynamics, free of most approximations that have plagued the modeling of these nuclear reactions for more than half a century. The quantum jumps from one potential surface to another, corresponding to various possible modes of decay and internal excitations of the fragments, will be described in a stochastic realization of the real-time path-integral formulation of the quantum many-body problem. In principle, this is an exact solution of the many-body Schrödinger equation, if the relevant part of the Fock space is explored.

This will be one of the main challenges—the implementation of a Metropolis importance sampling for the real-time quantum dynamics. The other major challenge will be the creation of a numerically stable and accurate propagation in time on an unprecedented scale. Without the element of quantum jumps, the nuclear dynamics reduces to a version of time-dependent DFT in the adiabatic approximation. In the case of the nuclear problem, this already is a petascale problem, which currently is under implementation and study within the SciDAC 2-UNEDF program (Bertsch et al. 2007). Allowing for the quantum jumps requires performing the same type of simulation numerous times (from thousands to millions of realizations) under stochastically generated conditions to capture all of the relevant quantum configurations of an excited fissioning nucleus. Not only would one be able to generate the partial cross sections for the large number of splitting modes of a fissioning nucleus, but also the cross sections describing much finer details, such as the excitation energy of the emerging fragments.

The long-term goal of all of these studies is to unify the fragmented knowledge of the nuclear fission process. At present, fission half-lives, fission cross sections, fission fragment properties, and prompt fission neutrons and γ-rays are all fission observables that are modeled (not even predicted in most cases) separately. A comprehensive nuclear fission theory and modeling should be able to predict all of those quantities consistently. In addition to understanding the processes, the uncertainties in the basic observables and the correlations between them must be carefully characterized. This is a particularly challenging task, both in setting up the correct framework to examine and characterize the uncertainties and in their evaluation. Characterizing uncertainties will typically require at least an order-of-magnitude increase in computational requirements (often even more) over the basic calculations.

The national security outcomes of a comprehensive treatment of nuclear fission are widespread and important. They affect not only traditional NNSA priority areas, but emerging concerns, including detection of nuclear materials, attribution, and energy production. By understanding possible fission signatures, particularly the time structure and detailed correlations of fission products, nuclear materials detection could be more effective. Advanced reactor concepts often require understanding fission in higher-energy regimes and for a wider arrangement of actinides than typically studied under the auspices of the NNSA. An improved theoretical description of fission aligns with the goals of the NNSA and its Stockpile Stewardship Program, which entails an accurate and complete modeling of the nation's aging nuclear weapons stockpile. All of these applications also require a sophisticated picture of the uncertainties in fission, further increasing the computational requirements (see Figure 8). Understanding fission is central to the continuing process of certifying both the safety and reliability of the stockpile without resuming nuclear testing and increasing the threat from nuclear proliferation.

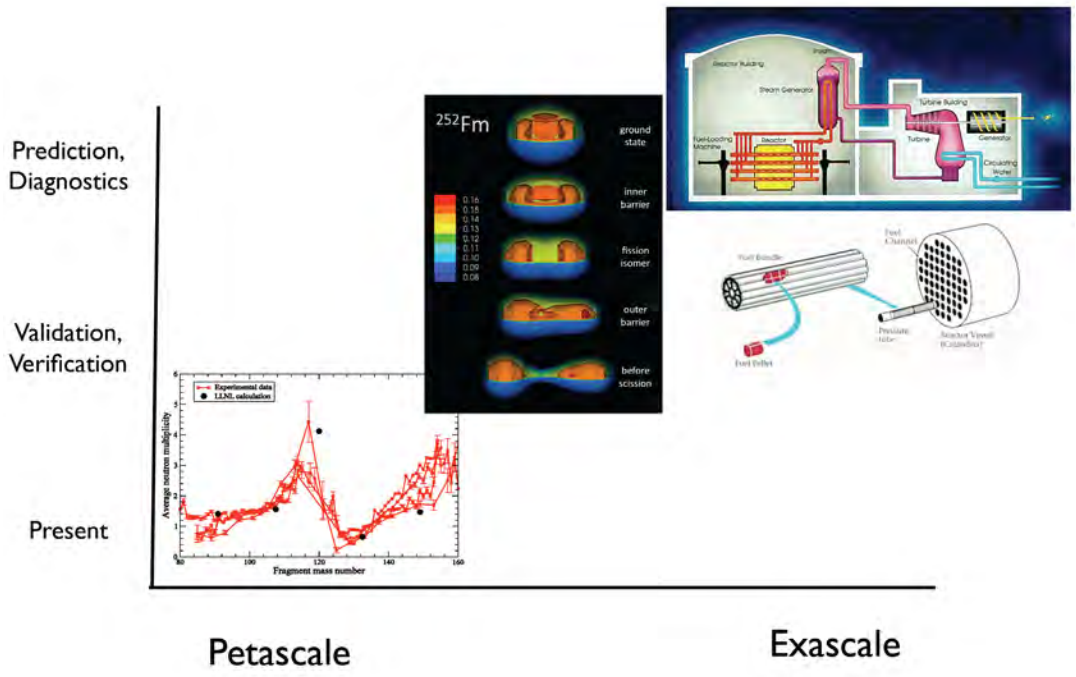

Figure 8. Computational advances in fission progressing toward exascale computing. Present calculations are predominantly fission barriers, although some calculations of neutron multiplicities (shown) are becoming available. Exascale computing will enable the study of fission dynamics and more predictive calculations of processes relevant to full-scale reactor simulations. Image courtesy of Joseph Carlson (Los Alamos National Laboratory) and Edward P. Hartouni (Lawrence Livermore National Laboratory).

Theory of Nuclear Reactions

Recent developments in theory and computation that improve the understanding of fission also will enable calculations of a variety of other reactions central to the NNSA program. Many of the most important reactions are neutron-induced, including neutron capture, (n,n') reactions, and (n,2n) reactions (Chadwick et al. 2007). The reactions will be performed with the highest physical realism starting from the correct fundamental description of the nuclear physics, and providing true predictive capability regarding the nuclear phenomena. Exascale computing will make these claims realizable.

Understanding and predicting the properties of energy-producing reactions of interest, as well as further reactions on the resulting fragments, will enable accurate simulations of nuclear systems. Presently, much of the data in the reaction libraries are obtained by fitting models to available experimental data. More sophisticated models will be required to understand the correlations between the relevant reaction rates and, in particular, how the uncertainties are correlated. The latter is particularly critical to obtain believable uncertainties at the system level (Chadwick et al. 2005).

Nuclear products that result from targets after interaction with neutrons incident with energies from thermal through a few MeV up to 150 MeV must be understood. All kinds of nuclear targets need to be studied, from light nuclei through fission fragments up to existing and newly produced actinides. The real problem is knowing this not just for isotopes that occur naturally and can be made into targets for neutron beams, but also for a wide range of radioactive nuclei that decay too quickly to be experimental subjects, including fission fragments, tracer elements, and a range of minor actinides. The radioactive nuclei cannot be made into beams incident on neutrons either because neutrons themselves decay too quickly.

These physics problems can only be solved with the help of a robust nuclear theory that has accurate predictive capabilities. Such theory can either be used *ab initio* to predict the needed reaction outcomes or to extrapolate from stable to unstable nuclei in some range of isotopes. There is an urgent need to develop a reaction theory that can be integrated with the best nuclear structure predictions. The new reaction theory will replace the older empirical models with a modern microphysics-based approach, and will at least properly deal with having several particles in the continuum because their correlations have never been treated accurately. The plan is to implement full a coupled-channels reaction theory with coherent exit channels using input transition densities and level densities from microscopic nuclear structure, including auxiliary-field shell models and random phase approximation.

New methods will require a fully quantum-mechanical treatment of all couplings and continuum states, a result well beyond what has been achieved with existing human labor and computer power just over teraflops. At present, there is only the capacity to generate and couple the states that connect directly to the ground state. Soon, including excited states at least two or three steps beyond those directly coupled will be necessary. To quantify uncertainties, it will be necessary to use Monte Carlo instantiations of the high-level densities of nuclear excited states, minimizing semi-classical approximations and seeking a full quantum-mechanical treatment of the production and decay of the intermediate compound nuclei for each instantiation. Until new human labor resources are available in the NNSA laboratories (at present in this field there are only five or six experts across the whole complex), the actual computational requirements cannot be accurately estimated.

Statistical approaches will continue to play an important role in understanding reactions, particularly for those involving compound nuclear states. Monte Carlo simulations of nuclear reactions are allowing an unprecedented level of detail in the representation of physical observables (Kawano et al. 2010). Moving beyond averages, these types of calculations can assess a vast array of correlated data, which can lead to significant advances in understanding nuclear structure properties and nuclear reaction mechanisms; for example:

- Now, (n,n'γ) reactions with a gate on a particular γ-ray can be assessed both experimentally (when statistics are large enough) and through modeling. This type of data can be used to infer nuclear-level density properties, such as spin and parity distributions in compound nuclei.

- Prompt fission neutrons and γ-rays emitted rapidly following the scission process can be used to constrain theoretical models of fission. However, stringent constraints can only be placed if correlated data and probability distribution functions are available. Averaged quantities are simply not enough.

In seeking finer and finer details and correlations in the set of calculated data, such Monte Carlo simulations rapidly can become heavy CPU consumers. The Monte Carlo simulation generates all correlated observables, such as correlations among neutrons and gamma-rays. This correlation information, which is difficult to obtain with a traditional deterministic methods, can be applied directly in radiation transport simulations by using the nuclear reaction model calculation as an internal Monte Carlo-event generator. In the case of rare events or correlations, the problem becomes acute and only intensive computing power, along with smart algorithms, would be able to provide answers.

Comprehensive and reliable understanding of neutron-induced reactions on all target nuclei, including unstable nuclei inaccessible in the laboratory, are critical for emerging NNSA priorities. Detection of nuclear materials requires understanding neutron, photon, and muon-induced reactions over a range of energies. To take maximum advantage of new detection schemes, the time- and space-correlations of reaction products are clear signatures that must be understood. Exascale computing will play an important role in developing these new capabilities.

Simulation of Integrated Nuclear Systems

The first three PRDs involve accurate calculations of the microscopic nuclear physics. For these advances to affect NNSA priorities, system-level simulations must advance to incorporate more sophisticated nuclear data and to more completely describe the integrated nuclear systems. Nuclear physics computations are critical to all nuclear security applications, including nonproliferation, attribution, and threat reduction. Bringing the nuclear physics to bear on important problems, including threat reduction and energy, also requires exascale computing.

The applications of interest are as follows:

- advanced nuclear energy systems, including not only reactor concepts, but all aspects of the fuel cycle

- global security applications, such as nuclear detection, nonproliferation, and nuclear forensics

- physics of nuclear weapons

- modeling and analyses of major, new above-ground experimental facilities and associated diagnostics, such as NIF and MaRIE

- safety applications, including nuclear criticality safety, radiation protection, and shielding

- medical applications, both diagnostic and therapeutic.

The common denominator in all of these applications is that simulation accuracy of the transport packages used in modeling critically depends on the accuracy of a large phase-space of nuclear data derived from nuclear physics theory, models, and experiments.

For example, the conceptual design of advanced reactors is expected to be increasingly based on high-fidelity modeling and simulations, while at the same time, the safety, construction, and economic margins of new nuclear power plants are becoming more stringent. Such a challenge requires large-scale computing to run detailed calculations of all nuclear reactions happening in a reactor core. In turn, those calculations have to be coupled with other physics computational tools in a multiphysics simulation code.

While the simulation of the regular operation (static) of a reactor may not need large-scale computing power, transient behaviors (such as the simulation of an accident) can only be simulated accurately with time-dependent (and CPU-intensive) methods. The propagation of uncertainties and correlations in the underlying physical properties of the reactor (nuclear data, material properties, etc.) also has to be taken into account rigorously, and usual multigroup approximations should no longer be used. The accurate design and simulation of an advanced reactor is only a piece of the overall nuclear fuel cycle, which includes the preparation of innovative fuels, the transmutation of nuclear waste, nonproliferation tools and techniques, etc. These issues all would greatly benefit from advanced transport calculations that only can be achieved through advanced nuclear reaction modeling and by using supercomputers.

Increasingly, large-scale simulations of all the systems described are being used in place of similarly large-scale mock-up experiments. Among the drivers for this paradigm shift are objectives of reduced cost, reduced time, reduced uncertainty, and increased physics discovery and understanding. The fidelity (and therefore, the reliability) of these simulations can be shown to have a direct correlation to the level of fundamental nuclear physics knowledge.

Introducing these advanced simulations to the exascale era will involve many of the same computational issues as the nuclear physics simulations themselves. For example, many of the large system simulations are based upon Monte Carlo techniques and branching random walks, as are some of the quantum simulations already described. Breaking down the random walks into localized domains and efficiently optimizing the walks across hundreds of thousands to millions of cores is a formidable technical task. Many of the techniques being studied to enable this type of simulation will be applicable to both the microphysics and system-level simulations.

Simulations are required for design or “normal conditions,” but also for upset or accident conditions. Another substantial driver is the need for scientific uncertainty quantification, which demands accurate knowledge of uncertainties (and correlations) of nuclear physics data. The latter leads to pursuing more fundamental understanding of nuclear physics through research efforts, such as those described in previous sections. Pushing these applications to realistic, full-system simulations, even with the present physics input, will require at least an order-of-magnitude increase in computation. Many simplifications

are made in current codes, particularly in the area of fission where knowledge of the correlation of neutron and fragment decay is typically included on only a rudimentary level.

Additional requirements for uncertainty evaluations and correct treatments of time- and space-correlations in nuclear processes are particularly critical for threat reduction, as well as other applications, and drive the need for exascale computing. In addition to a better understanding of the microscopic nuclear physics, improved algorithms and exascale computers are required to retain complete fidelity in the simulation. The goal is to have tools that are predictive and reliable enough to make critical decisions in nuclear security.

CONCLUSIONS

Exascale computing will bring dramatic new capabilities to theoretical nuclear physics and related fields. These advances will enable accurate microscopic calculations of fusion reactions, vastly improved models of nuclear fission and neutron-induced reactions, and high-fidelity incorporation of these data and uncertainties into simulations of integrated nuclear systems. These advances will enhance traditional NNSA applications, including stockpile stewardship, and dramatically advance new NNSA priorities, including nuclear material detection, attribution, and energy security.

Initial calculations of fusion reactions have been undertaken with petascale resources and are quite successful. Extending these algorithms and techniques to more complex reactions will challenge the theoretical techniques and computational resources of next-generation facilities. To predict unmeasured cross sections and accurately include the effects of multi-nucleus final states will require exascale computational resources. These advances will further research in basic and applied nuclear physics, affecting fields such as big-bang nucleosynthesis and stellar evolution, as well as NNSA priorities. Advanced calculations of fusion reactions will enable enhanced diagnostics of NIF capsules and other environments where fusion plays an important role.

Fission is a long-standing NNSA priority, and computational efforts are having an important impact. Present-day efforts are improving DFTs of nuclei, concentrating on fission barriers and related quantities. Future advances will treat the dynamics of fission, requiring exascale computational resources. Advanced models of nuclear fission will allow reliable treatment of correlations among the many types of nuclear fission data (fragment mass distributions, time- and space-correlations of neutron emission, etc.). Sophisticated models of nuclear data uncertainties will also be created. These new capabilities are critical for the detection of nuclear materials and modeling advanced reactor applications.

Exascale computing will also enable simultaneous advances in understanding a host of neutron-induced and related reactions. For example, sophisticated coupled-channel techniques will enable treatment of unstable nuclei reactions that are extraordinarily difficult to measure experimentally.

Understanding the uncertainties in reaction models, combined with large-scale simulations of integrated nuclear systems, will provide accurate, comprehensive pictures of systems, spanning NIF capsules, active and passive interrogation of nuclear materials, and advanced reactor concepts.

Exascale computing and corresponding theoretical advances will afford entirely new capabilities within nuclear science. Computational nuclear science will proceed congruently with experimental science as a critical part of understanding nuclei and their reactions, particularly in regimes where experiments are

difficult to impossible. Exascale computing will be essential to make this new knowledge work effectively toward achieving important NNSA priorities in global security.

Table 1 lists the milestones for the work described in this section. Provided that the computational resources become available for research at the anticipated scales, the forefront research activities that will be conducted are provided as milestones.

Table 1. Milestones for Nuclear Physics

Scale	Milestone
>1 Petaflop-year	• Nuclear Physics – Fusion of lightest nuclei (dt) and one- and two-channel scattering/reactions – Characterization of fission potential energy surfaces and static properties – Calculations of nuclear reactions including (n,2n) etc. with greater resolution
>20 Petaflop-years	• Nuclear Physics – Fusion of larger nuclei (A<12) with two-nucleus states – Initial treatment of fission dynamics: determination of inertia tensor – Accurate studies of (n,n') for nuclear-level densities
>100 Petaflop-years	• Nuclear Physics – Fusion with multi-nucleus final states: ^{6}Li (t,n)2 ^{4}He and ^{7}Li(d,n) 2 ^{4}He – Fission dynamics with time-dependent generator coordinator method – Characterization of prompt fission neutrons and gamma rays for global security – Realistic NN phase shifts from lattice QCD
>1 Exaflop-year	• Nuclear Physics – Neutron-induced fusion at higher energy for NIF diagnostics – *Ab initio* treatments of screening in nuclear reactions – Induced fission with Monte Carlo evaluation of path integral for real-time properties – Realistic three-nucleon force (including three neutrons) from lattice QCD

MATERIALS SCIENCE

Co-Leads: John Sarrao, Los Alamos National Laboratory
Sidney Yip, Massachusetts Institute of Technology

Panel Members: Ron Ballinger, Massachusetts Institute of Technology; Corbett Battaile, Sandia National Laboratories; Jim Belak, Lawrence Livermore National Laboratory; Vasily Bulatov, Lawrence Livermore National Laboratory; Alfredo Caro, Lawrence Livermore National Laboratory; Michael Demkowicz, Massachusetts Institute of Technology; Stephen Foiles, Sandia National Laboratories ; Timothy Germann, Los Alamos National Laboratory; Yogi Gupta, Washington State University; Christian Mailhiot, Lawrence Livermore National Laboratory; Eric Schwegler, Lawrence Livermore National Laboratory; Marius Stan, Los Alamos National Laboratory; Priya Vashista, University of Southern California; Arthur Voter, Los Alamos National Laboratory; John Wills, Los Alamos National Laboratory; and Robert Averback, University of Illinois at Urbana-Champaign

CURRENT STATUS

In the last decades, the National Nuclear Security Administration's (NNSA) national security challenge has evolved, and the role of simulation and computation has grown dramatically. The process of certifying nuclear weapons performance has changed from one based on integrated tests to science-based certification in which underground nuclear tests have been replaced by large-scale simulations that are appropriately validated with fundamental experimental data. Further, the breadth of national security challenges has expanded beyond stewardship of a nuclear deterrent to a broad range of global and asymmetric threats.

Materials challenges are central to the full suite of these national security challenges. Mission requirements demand that materials perform predictably in extreme environments such as high pressure, high strain rate, and hostile irradiation and chemical conditions. While considerable advances have been made in incorporating fundamental materials physics into integrated codes used for component certification, significant uncertainties still remain—and materials properties, especially at the mesoscale, are key to understanding uncertainties that remain in integrated weapons performance codes and that are currently treated as empirical knobs. Further, additional national security mission challenges could be addressed more robustly with new, higher-performing materials.

Current models of materials strength, and the onset of damage and eventually failure, cannot adequately account for mesoscale heterogeneity intrinsic to real materials. For example, during dynamic compression, materials properties such as phase, density, temperature, and the onset of fracture would benefit from improved models at the crystal-to-grain scale. A number of these phenomena are intrinsically path- and time-dependent, due to dissipative effects and intrinsic inhomogeneities, and those characteristics place formidable demands on predictive models. These challenges are particularly acute when scientists must rely on models to predict materials behavior in experimentally inaccessible extremes that affect weapons performance.

Materials compatibility and aging effects also represent a significant challenge. The effects of corrosion and other interfacial chemistry phenomena are poorly predicted in many materials of interest. Similarly, aging effects, often but not exclusively due to radiation damage and radiolytic processes, affect materials

performance and lifetime prediction. Most acute, in fact, are the coupled consequences of irradiation and corrosion. The ability to predict materials performance and mitigate negative consequences must span both the decadal scales of warhead lifetime and the very short time scales of nuclear performance.

Finally, materials properties depend on the microstructure of the material, which is a direct consequence of processing. However, present materials specifications are process based rather than product based. Scientists rely on making a given material exactly as they have in the past for confidence in its performance. In many cases, this is no longer possible, due to new regulatory restrictions or the loss of availability of commercial vendors. In such instances, scientists' ability to confidently predict form, fit, and function equivalency is compromised. Product-based certification, in which confidence derives from a predictive understanding of materials functionality, promises a faster, cheaper, and better route to materials certification. In an era of a decrepit nuclear weapons production infrastructure that needs multibillion-dollar reinvestment, progress in predictive materials performance can pay enormous financial dividends.

The advent of exascale computing provides an opportunity to advance the frontiers of scientific understanding of materials performance, including the effects of materials fabrication and processing on performance (and materials failure)—i.e., process-aware performance. If achieved, such predictive capabilities can transform scientists' manipulation of microstructure and interfaces, especially at the grain scale, to enable the design and development of extreme-environment-tolerant advanced materials. The scale (in spatial and temporal dimensions) of simulation enabled for the first time with exascale computational resources is exactly equivalent to the mesoscale of materials from which the greatest current uncertainty derives.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

The materials challenges that need to be addressed to meet these national security mission needs are well documented (DOE 2006; DOE 2008). Other workshops in this DOE Office of Advanced Scientific Computing Research series have also addressed the opportunity for exascale computing for broadly defined materials challenges (e.g., Scientific Grand Challenges through Extreme Scale Computational Science for Nuclear Energy and for Basic Energy Sciences[5]). In all cases, advances in computational power coupled with advanced theoretical models, and validated by experiments, provide great promise in achieving a science-based, process-aware understanding of materials performance.

The overarching grand challenge is to achieve a predictive, mechanistic understanding of real materials. In these circumstances, microstructures and interfaces matter, kinetics matter, and environments matter. With the advent of exascale computing, the possibility now exists to achieve predictive capabilities to manipulate microstructures and interfaces, at the grand scale, to enable the design and development of extreme-environment-tolerant advanced materials. The exascale challenge, in essence, is direct access to materials functionality at the mesoscale frontier: predicting mechanisms governing materials properties in micron-scale volumes over times of milliseconds. While framed in this panel report as specific to NNSA, these issues are a microcosm of broader materials challenges, especially in extremes such as materials needs for nuclear energy.

[5] http://www.sc.doe.gov/ascr/Misc/GrandChallenges.html.

The Exascale Challenge is the Mesoscale Frontier

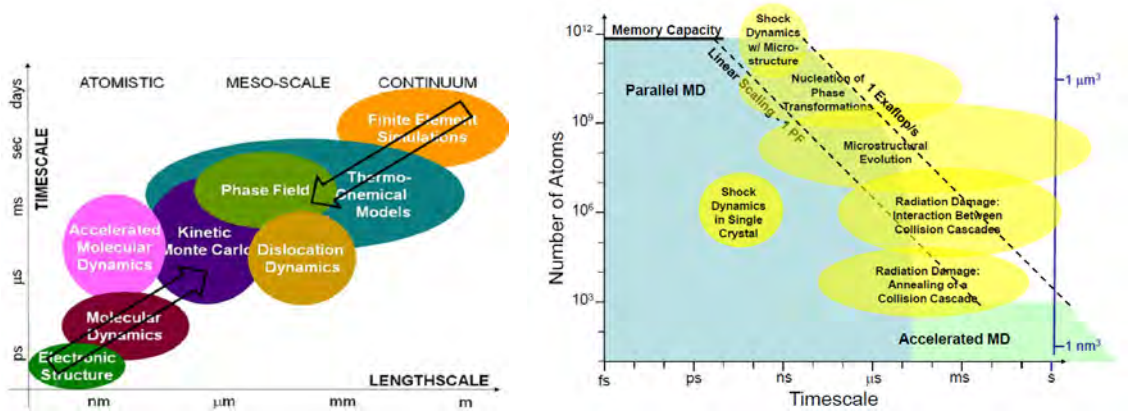

The grand challenge of multiscale modeling is to connect atomic scale processes to bulk continuum models. The current state of the art is the effective passing of input and output parameters between various models and approaches (left figure). However, the challenge is not to run atomistic models on the meter scale. Rather, it is to connect these extremes at the intermediate mesoscale and to use such mesoscale models to establish requirements for atomic scale calculations that then directly affect continuum approaches. As the right figure indicates, many of the mesoscale materials properties that researchers need to understand require exascale computational resources to span relevant temporal and spatial scales. Images courtesy of John Sarrao (Los Alamos National Laboratory) and Sidney Yip (Massachusetts Institute of Technology).

Success in this endeavor will require strong coupling and feedback among experiment, theory, and computation. A number of experimental capabilities—from the Linac Coherent Light Source at Stanford Linear Accelerator Center to the National Ignition Facility at Lawrence Livermore National Laboratory for example—are presently coming online and will allow access to previously unprecedented extreme environments and unprecedented spatial and temporal measurement resolution with which to measure materials properties. Additional experimental facilities are being planned that would enable unprecedented control of materials functionality through in situ, real-time measurements. Thus, the co-design of exascale computational resources is very timely to promote the necessary synergy between theory, experiment, and computation.

PRIORITY RESEARCH DIRECTIONS

To address the overall grand challenge of achieving a predictive, mechanistic understanding of real materials, this panel identified four specific priority research directions (PRDs):

- mechanistic understanding of environmental degradation, coupling chemistry and stress
- atomistic simulation on engineering time scales
- predict constitutive representation in multiphase materials
- electronic structure theory and simulation methods for nonadiabatic and strongly correlated systems.

The first and third PRDs, which are particularly acute for national security mission requirements, directly address the challenges in materials durability and materials strength and damage described previously. They are examples of "mission pull." The second and fourth PRDs are focused on the development of advanced methods and algorithms; they typify "science push." Success in these two PRDs will directly affect the application areas of the other two PRDs.

Significant overlap also exists between these PRDs and those of the Chemistry panel, especially relative to the need for enhanced electronic structure and dynamics methods. The focus on interfaces implicit in PRD 1 is central to surface chemistry; however, the emphasis on the interaction between chemical and mechanical effects is differentiating. Similarly, the relatively long time and length scales associated with mesoscale materials interactions extends the electronic focus of chemistry to lattice degrees of freedom. Better performing and safer high explosives are a key challenge for nuclear weapons that spans both chemistry and materials science.

Finally, each of these PRDs represents a computational challenge that is intrinsically exascale in nature and cannot be addressed simply by the performance of many petascale calculations. Exascale computation would allow direct atomistic simulations of trillions of atoms for microseconds; the extension of intrinsically tightly coupled dynamic processes to millisecond time scales; simulations with realistic defect densities and their evolution for nanoseconds; and the simulation of, for example, melt curves for strongly correlated systems like plutonium for the first time.

Interfacial Chemomechanics

Scientific and Computational Challenge

The challenge is to gain a mechanistic understanding of environmental degradation in which chemistry and stress effects are emphasized equally through a combination of electronic structure, atomistic, and mesoscale simulations and validation experiments. Corrosion is a ubiquitous phenomenon central to many scientific studies and technological applications. It could be the oxidation of a high-temperature ceramic component in an aircraft engine or the hydrolytic weakening of quartz in the Earth's mantle; weapons materials such as uranium are also known to corrode in the presence of oxygen, hydrogen, and water vapor. Regardless of the specifics, scientists can identify a fundamental process that is vital to understanding the chemomechanics of materials. At the conceptual level, the most significant questions to address are the interactions between the environment, the surface film, and the underlying substrate; the presence of an embrittling (corrosive) species; and the spatial-temporal evolution of the film-substrate system to the point of "failure." In this context, a phenomenon of very broad interest is stress-corrosion cracking (Ciccotti 2009). One attempt to approach this process through multiscale modeling is to focus on the formation of a passive ultrathin oxide film and its evolution to the onset of structural breakdown (Yildiz et al. 2009). At the molecular level, scientists need to treat charged defect transport and electron transfer in the aggregation of cation vacancies that eventually lead to film-substrate decoherence and pit nucleation. This is an example of combining unit process (charged transport) to study system-level behavior (film growth). Furthermore, by considering substrates of iron and nickel-chromium-iron alloys in a systematic study of pit nucleation and crack propagation simulations (under controlled ambient and high temperature and stress environments), scientists can hope to achieve predictive insights into transgranular versus intergranular cracking in structural materials.

Summary of Research Directions

Research in interfacial chemomechanics is recommended to accomplish the following:

- extract mechanisms from reactive molecular dynamics directly linking unit processes to integral behavior over microns and milliseconds

- develop robust, portable, O(N) codes for multicore architectures building on existing petascale codes.

Research directions for the study of stress corrosion cracking (SCC), as discussed in the following sidebar, would consist of setting up a three-dimensional model of a crack tip to investigate the initiation effect of a corrosive agent. After the basic simulation model has been formulated and thoroughly studied, scientists can then vary the stress intensity factor to directly study the behavior in regions 2 and 3 addressed in the "Stress Corrosion Cracking" sidebar. To probe transport processes on long time scales (crack growth is known to be a slow phenomenon, beyond the reach of current time-scale capabilities of molecular dynamics), scientists would need to develop simulations methods that can reach much longer times. This is a highly nontrivial challenge (see PRD, "Atomistic Simulation on Engineering Time Scales").

A specific problem for further research could be the interaction between water and silica, a relatively simple system, where the bond-strain effect on the reaction H_2O + Si-O-Si→ 2SiOH has been characterized (Zhu et al. 2005). Considerable computational resources are in place to study the system of amorphous silica and water (Nomura et al. 2009), including billion-atom molecular dynamics and microsecond molecular-dynamics simulations of damage initiation and growth mechanisms during fracture, deformation and plasticity of voids under shear, and defect nucleation and migration during nanoindentation (Chen et al. 2009).

To advance to the next stage in the coming era of exaflop computers based on clusters of many-core processors, it is essential to develop a scalable parallel and distributed computational framework consisting of methods, algorithms, and integrated software tools for SCC simulations with quantum-level accuracy (see figure in "Stress Corrosion Cracking" sidebar). Specifically, researchers should develop the following:

- metascalable (or "design once, scale on new architectures") simulation framework (Nomura et al. 2009) in conjunction with linear-scaling algorithms and an auto-tunable parallelization framework for million-atom density functional theory, multibillion-atom reactive molecular dynamics, and trillion-atom and molecular-dynamics simulations

- automated dynamic model transitioning to embed high-accuracy simulations in coarse simulations on demand

- billion-atom molecular dynamics combined with accelerated dynamics schemes to reach microsecond time scales

- visualization and analysis of exascale simulation data.

Stress Corrosion Cracking

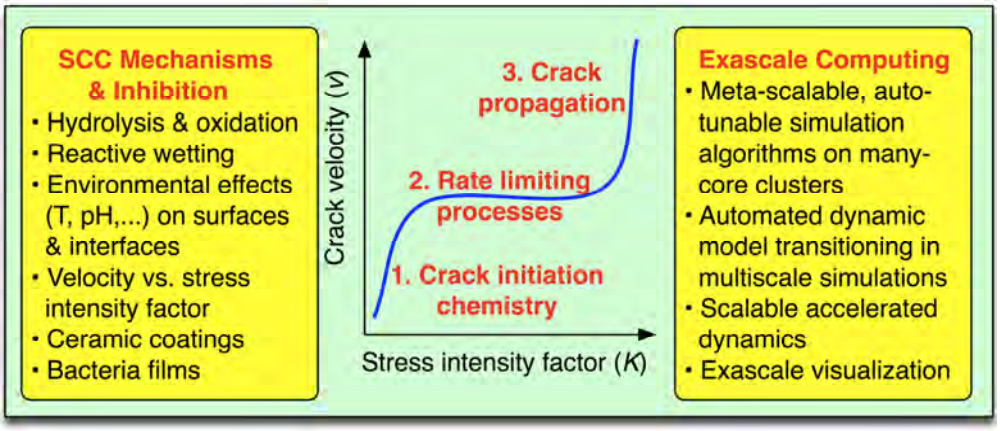

This figure depicts a crack moving at a steady speed v in a corrosive medium under stress (stress intensity factor K). The variation of the crack-tip speed v (usually shown as log v) with applied stress illustrates the classic behavior of stress corrosion cracking (SCC), described in qualitative terms through a variety of phenomenological models and commonly used in data interpretation. Image courtesy of P. Vashishta, R. Kalia, and A. Nakano (University of Southern California).

As illustrated in the above figure, the onset of crack propagation at low stress (region 1) suggests a threshold effect where bonds break at the crack tip primarily due to corrosion attack. The rapid increase of $v(K)$, essentially exponential, demonstrates a strong stress dependence as well as sensitivity to the environment. In region 2, crack speed becomes stress independent because transport of a corrosive agent is the limiting process. Modeling this region properly requires a full treatment of the interplay between local reactivity and spatially extended transport, which endows in SCC its hallmark characteristic of slow dynamics. In region 3, the propagation curve reverts back to a second threshold behavior: brittle fracture with corrosion essentially playing no role.

The behavior of the constitutive response curve—two regions of different governing mechanisms of crack-tip extension separated by an incubation range—seems to be a generic characteristic of chemomechanics problems. For example, the well-known hydration kinetics or setting curve of cement has a corresponding form to the sidebar figure. In this case, the shear modulus of the cement paste has a time dependence showing an initial increase to MPa range (the coagulation stage, region 1), followed by a plateau (incubation stage, region 2), and the onset of hardening where the modulus increases to the GPa range (setting stage, region 3) (Lootens et al. 2004). Despite this qualitative level of understanding, no quantitative models at the molecular level exist that can capture the overall characteristic behavior of SCC (Ciccotti 2009). The reason is that the chemomechanical processes governing this curve require modeling studies and computational capabilities at a level of sophistication that are only beginning to be developed. The advent of exascale computing brings the study of corrosion and crack extension into the range of appropriate research objectives.

Such an exascale hierarchical multiscale computational framework will enable simulations of SCC in metals, alloys, glasses, and high-temperature structural ceramics in gaseous and aqueous environments (see figure in “Stress Corrosion Cracking” sidebar). Simulations will provide an atomistic-level description from which researchers will extract the basic mechanisms of SCC. This understanding will

feed into atomistically informed continuum models to make contact with engineering SCC data for purposes of validation. It will also be possible to investigate SCC inhibition by ceramic coatings (e.g., alumina and silicon carbide), bacteria films, and self-assembled monolayers (e.g., oleic imidazolines).

Returning to the general challenge of environment-surface, film-substrate interactions, in all but the most aggressive environments the cracking process consists of an incubation period where the conditions are established for crack initiation, followed by crack initiation and subsequent propagation. This process is illustrated by the schematic in Figure 9. In most cases stages I-III dominate the life of the component. During these stages, various species in the environment play key roles in both the establishment of the surface film and the local destabilization of the film to allow entry of the embrittling species. Embrittling agents include oxygen, hydrogen, and other species such as halogens (Cl^-, sulfate, etc.) that act to enhance the effectiveness of the process. A key to elucidating this interaction will be to use modeling and simulation, coupled with high-fidelity experimental investigations, to gain an understanding of the nature of water, anion, and cation interactions on the surface and their interaction with the metal atoms including the effect of stress/strain.

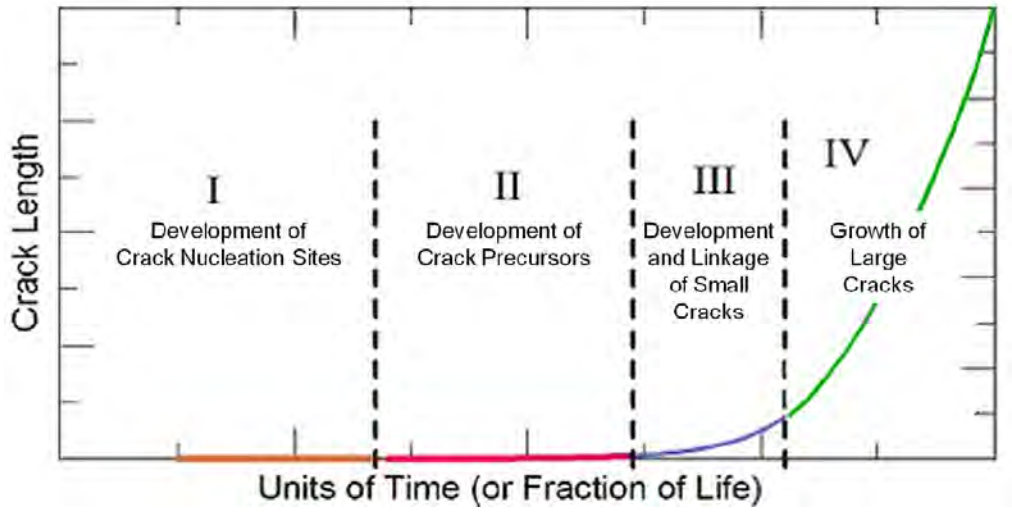

Figure 9. The evolution of the cracking process from crack initiation to subsequent propagation over the lifetime of a crack (Pathania 2005). Interactions with the environment play a strong role in crack evolution. Image courtesy of R. Parthania (Electric Power Research Institute).

At the interface between the film and the substrate, the action of the embrittling agent to gain entry to the material is another critical area. Often this involves the grain boundary or other higher energy two-dimensional surface (twin boundary, etc.). It is known that the nature of the grain boundary (degree of coincidence with the lattice) plays an important role in the process. It is also known that the type of interaction between the embrittling agent and the grain boundary depends not only on the degree of coincidence but also on what is on the grain boundary. For example, the presence or absence of oxidizable or reducible species/conditions with respect to the bulk alloy chemistry, contaminants, and precipitates is a critical factor in the process. Recent work has shown that while oxygen plays a key role, if the source of oxygen is the dissociation of water, hydrogen also will be present. The development of an understanding of processes in this region will require the combination of high-fidelity modeling and simulation with high-resolution experimental work.

The third element in the environmental degradation "triad" is to understand how the embrittling agent acts to promote crack extension. The crack extends by the breaking of atomic bonds. While there is controversy over the exact nature of the process, the combination of high-fidelity modeling and simulation and experimental investigation promises to finally shed light on this problem. In this region, the interaction between the embrittling agent and an evolving chemistry presents unique problems and opportunities for research. One can consider the actual degradation process to be an initiation followed by crack propagation in a sequence of such events. The evolution of the grain boundary chemistry due to species migration along the grain boundary to a "process zone," as well as from within the matrix, results in a time-varying chemical potential in the process zone. Add to this the presence of stress/strain and radiation damage, and the complication multiplies. Understanding how this environment evolves dynamically with time/dose is essential to understanding how the embrittling agent will react. For the case where oxygen is the embrittling agent, several interactions are possible: direct oxidation of primary alloy element species; oxidation of contaminants; or oxidation reaction with precipitates. The reaction process itself will create a chemical potential gradient that will promote mass transport. The solution to this problem will, as with the other regions, require the combination of high-fidelity modeling and simulation with high-resolution experimental techniques.

Expected Computational and Scientific Outcome

The use of multiscale modeling concepts and simulation techniques to study the destabilization of an ultrathin layer of oxide interface between a metal substrate and the surrounding environment is an example of the notion of concept materials (Yip 2009) in the emergence of computational materials (Yip 2003). In this case, issues range from dealing with unit processes in mechanical deformation and crack propagation to treating chemical reactivity and charged-defect transport in an integrated fashion. These give rise to a set of interrelated challenges that include understanding electron transfer processes using first-principles methods; modeling cation transport and associated charged defect migration kinetics; and simulation of pit nucleation and intergranular deformation to initiate the breakdown of the oxide interlayer. These challenges also need to be addressed in an integrated manner at the system level. Such a problem illustrates a level of multiscale complexity that would be intractable by other means; it also points to a perspective framework that could guide future research in the broad community of computational science.

Continuing developments in computing technologies are providing unprecedented opportunities for the simulation study of SCC beyond recent work (Zhu et al. 2005; Ciccotti 2009). State-of-the-art atomistic simulations include 218 billion-atom molecular-dynamics simulation and 1.68 trillion electronic degrees-of-freedom quantum-mechanical calculation in the framework of density functional theory, with parallel efficiency over 0.95 on 212,992 BlueGene/L processors (Nomura et al. 2009). The spatiotemporal scale covered by molecular-dynamics simulation on a sustained petaflops computer per day is $NT = 2.1$ (e.g., $N = 2.1$ million atoms for $T = 1$ microseconds) (Nomura et al. 2009), beyond which a variety of accelerated-dynamics simulation methods are being developed (Vashishta et al. 2008). Scientists have also performed multiscale quantum-mechanical/molecular-dynamics simulation on a global grid of supercomputers, in which quantum-mechanical calculations were adaptively invoked only when and where high accuracy was required (Takemiya et al. 2006). Furthermore, to seamlessly couple quantum-mechanical and molecular-dynamics simulations, an intermediate layer has been introduced in which first-principles-based reactive force fields were used to describe chemical reactions involving more than one billion atoms (Nakano et al. 2008).

Potential Impact on National Security

Materials failure at high temperatures and stresses, and in harsh chemical and radiation environments, is a timely scientific challenge with far-reaching impact in major national security and nuclear technology enterprises. It involves premature and catastrophic failure as a result of a complex combination of stresses and corrosive reactions further accelerated in the presence of high radiation fluxes. Atomistic-level understanding of corrosion initiation mechanisms is critical to microstructure optimization leading to the design of materials resistant to environmental degradation in general and stress corrosion cracking specifically. This is possible through the development of chemomechanical models capable of describing crack tip behavior in an environment of high temperatures, chemical attacks, and concentrated stress loading. Such models will be necessarily more sophisticated than the crack propagation models currently available. Using atomistic simulations to demonstrate stress corrosion cracking, different size domains will have to be considered with each requiring a different computational approach. The domains must be nested because regions far from process zones do not require the detailed description necessary for accurate prediction near the crack tip. Experience from these simulations will be useful for extracting atomistic-level understanding of a broad class of materials-failure phenomena where corrosion plays a significant role.

Significant impact on national security in the arena of materials science and technology will also stem from the ability to design materials properties specific to their performance in contrast to the more traditional empirical approaches. Figure 10 depicts three examples to illustrate the potential of microstructure-based tailoring of interfaces in national security-relevant materials pertaining to SCC (Hosemann et al. 2008), mechanochemistry of foams (Li and Misra 2009), and nanostructured materials (Mara and Rollett 2009).

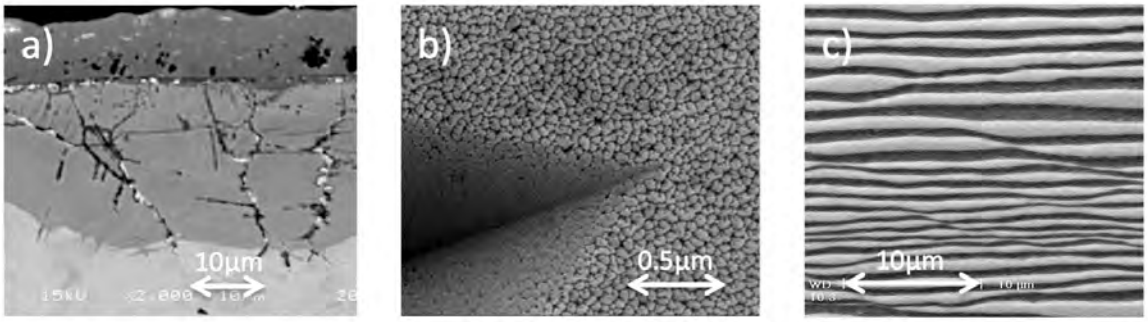

Figure 10. Micron-level physics governs a) stress corrosion cracking, b) mechanochemistry of nanofoams, and c) interface-dominated nanocomposites. Image courtesy of Michael Demkowicz (Massachusetts Institute of Technology).

The collective operation of unit processes manifests itself through phenomena like interaction of dislocation cell walls with heterophase interfaces, precipitate structure evolution during diffusion of impurities along grain boundaries, or pitting of oxide layers in corrosive environments—all of which take place over micrometers and milliseconds, or longer. Current computers are not amenable to atomistic modeling studies over this range of lengths and times.

Building on existing codes designed for petaflop computers (Nakano et al. 2008; Germann and Kadau 2008), exascale machines will, for the first time, make the micrometer/millisecond regime of materials

response accessible to direct molecular-dynamics simulation. This capability will make possible systematic validation and improvement of higher length-scale modeling tools such as dislocation dynamics or phase field methods by direct comparison with simulations performed with atomistic fidelity. The emergence of such large-scale molecular-dynamics simulations will also mark the first time that atomic-level predictions of the behavior of complex heterogeneous materials can routinely investigate phenomena on the same spatial and temporal scales as science-driven experiments, which are also beginning to approach micron- and millisecond-level resolution (DOE 2009).

Atomistic Simulation on Engineering Time Scales

Scientific and Computational Challenge

Accurate multiscale models for microstructural evolution on engineering time scales and deformation and failure under extreme conditions or environmental attack are needed to meet certification challenges for nuclear weapons. Materials challenges crucial for national security span a wide range of issues (e.g., radiation damage, thermal and irradiation creep, corrosion, plastic deformation, and material failure). To address these phenomena, modeling methods that can reach time scales of milliseconds and beyond, as well as length scales of microns and beyond, need to be developed. A significant slice of essential physics lies at the mesoscopic scale, involving processes such as grain growth, annealing and interaction of radiation damage cascades, nucleation of phase transformations, and entanglement and depinning of dislocations. Multiscale approaches such as dislocation dynamics, thermochemical modeling, and phase field models can be applied in this regime—but they need to be carefully parameterized and benchmarked. Molecular dynamics, with its atomistic fidelity, can provide the necessary input and validation results for the higher-level methods. However, molecular dynamics, in its traditional implementation, has difficulties with reaching long time scales. This is a longstanding challenge that is still mostly unresolved.

The advance to the exascale will be dominated by increased concurrency, not increased processor speed. While length scales are easily extended as computers become more parallel (e.g., quadrillion-atom simulations should be possible on exaflop machines), time scales are not by virtue of their sequential nature. With a molecular-dynamics code parallelized through spatial decomposition, space can be traded for time without loss of efficiency only if the number of atoms on each processor remains large (e.g., 10^3 or more, although this is architecture dependent) so that the communication overhead required between force calls does not dominate. The result is that direct molecular-dynamics simulations beyond microseconds may still be exceedingly difficult, even on exascale machines. Accelerated molecular-dynamics methods, which operate by reducing the waiting time between infrequent events, can vastly improve on this, but currently they tend to be limited to small systems (10^3 to 10^4 atoms). Moreover, the computational speedup they can offer is severely hampered when the infrequent events are dominated by low barriers, as often happens in realistic systems. The result is there remains a substantial and important region of the length-time phase space that researchers could not access if they were given an exascale computer today.

Another aspect of the time-scale challenge is the existence of a wide spectrum of time scales associated with materials simulation. Unfortunately, microstructure evolution often involves mechanisms and processes that populate the entire spectrum of relevant time scales, from picoseconds to years. Damage accumulation in materials subjected to high-energy particle irradiation is an example where researchers must resort to rate theory approaches in the absence of suitable capabilities for long time simulations. In

addition, the challenge of long time scales manifests itself as a rare event problem where most of the computing cycles are spent on repetitive but uninteresting events (such as thermal vibrations in molecular dynamics) that rarely lead to a significant event such as diffusive rearrangement of atoms. Such rare events can themselves be unit steps eventually leading to even less frequent events such as vacancy detrapping from an impurity atom or cluster. The hierarchy of rare events may not end there and may include rare events of still higher order (e.g., formation of a supercritical nucleus of phase II). To gain access to relevant (long) time scales, it is desirable to somehow skip (or integrate over) the uninteresting unit events and focus computational effort on the interesting rare events. Rare-event sampling is a vast area of current research not only in material sciences, but also in chemistry, biology, and systems engineering. However, it is uncertain whether efficient methods for rare-event sampling on exascale computers can be developed.

An illustration of the capabilities and restrictions of molecular dynamics is provided by a recent simulation of grain growth in nickel, as seen in Figure 11. The results demonstrate annealing in nanocrystalline systems produces grains with incorporated defects, most notably twin boundaries in this case. This is consistent with experimental transmission electron microscopy observations of abnormal growth in annealed nanocrystalline nickel films (Hattar et al. 2008), while at conventional scales growth produces largely defect-free grains. The shortcomings of such calculations are related to nanoscale systems and high temperatures. The former has manifestations in both time and length scales, while the latter is necessitated by limitations to short times of simulation. For these reasons, it is not currently feasible to directly study the transition from growth of defected grains to the growth of defect-free grains. One could resort to mesoscale models that treat the motion of grain boundaries. Recent molecular-dynamics work has tried to quantify the required boundary properties, mobility and energy, for the full range of boundary crystallography (Olmsted et al. 2009a, 2009b). By construction, this class of models does not include the formation of defects during the growth. For this reason, the atomic-scale resolution needs to be retained. The availability of increased computing power (i.e., the exascale) would help but will not be sufficient. What is required is the application of an accelerated molecular-dynamics methodology. This will require significant advances on the current accelerated molecular-dynamics methodologies. First, this is a problem with a wide range of barriers and thus a wide range of rates. Second, this requires the treatment of orders of magnitude more atoms than have been treated with existing accelerated molecular-dynamics methods. Thus, such research areas would benefit from the implementation of accelerated molecular dynamics on exascale computers and of significant advances in the accelerated molecular-dynamics methodologies.

Fidelity and computability will remain two principal challenges for materials simulations on exascale platforms of the future. Much attention in the materials modeling community is currently focused on improving the accuracy of material models whereas just as much (if not more) progress can be achieved through improvements in simulation efficiency. To quantify, by simulation, the relationship between synthesis, processing, microstructure evolution, and resulting material properties a simulation should cover relevant (often large) time and length scales. Of the two scale challenges—time and length scales—the former is less likely to benefit directly from running existing material simulations on ultra-parallel exascale platforms. This is because space can be effectively partitioned whereas time cannot (time integration is inherently sequential). Thus, extending time scales accessible for material simulations remains a difficult challenge and demands most serious attention.

It is generally recognized that transport and deformation processes can be well studied using molecular dynamics only on the time scales of nanoseconds to perhaps microseconds. Phenomena occurring on

time scales of seconds or longer, relevant to the understanding of creep and related slow structural relaxations, are traditionally out of reach for atomistics unless an appropriate acceleration is available. Among the alternatives proposed, accelerated molecular-dynamics methods are the best understood and most extensively developed. The kinetic Monte Carlo (KMC) method, as a general class of stochastic methods, can be used to reach macroscopic times. An energy-landscape method, known as metadynamics, is relatively recent. Each method has been applied to a different problem to demonstrate its efficacy. Alternatively, there has been little or no intercomparison. Given the importance of time scale bridging in many of the materials simulation challenges, it would be useful to know more details about the advantage of one method relative to the others.

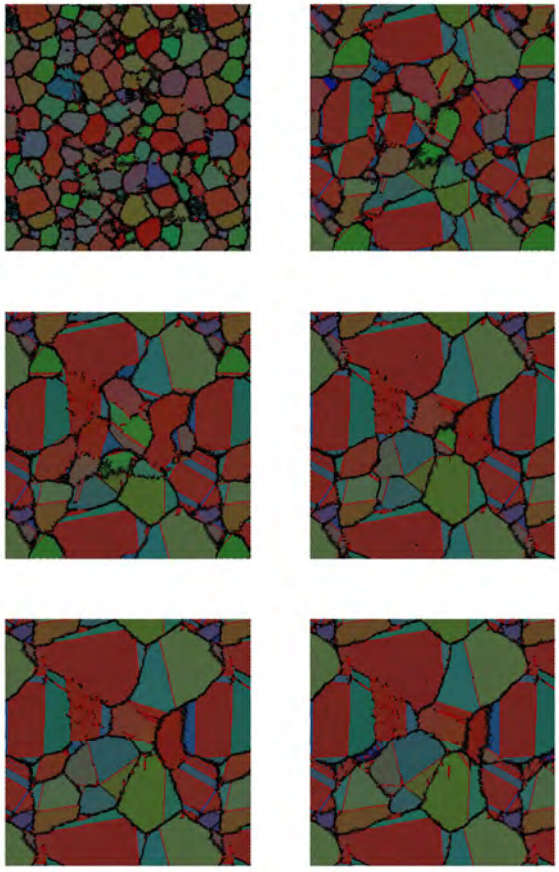

Figure 11. Representative snapshots of molecular-dynamics-simulated grain growth. These are snapshots of a representative slice through the simulation cell for T/TM = 0.75 at times of 0 to 5 nanoseconds in 1-nanosecond increments. The typical grain diameter in the initial state is 5 nm. The black atoms are non-crystalline atoms; the red atoms have hexagonal close packed nearest neighbor environments. The other colors have face-centered cubic nearest neighbor environments with the shade representing the local crystal orientation. The images show a region larger than the periodic boundaries of the computational cell. Image courtesy of S. M. Foiles (Sandia National Laboratory).

Summary of Research Directions

Research in atomistic simulation on engineering time scales is recommended to accomplish the following:

- Develop methods and algorithms to fully exploit extreme-scale computing for accurate long-time molecular-dynamics simulations of materials processes.

- Develop methods for optimal space-time partitioning on extremely parallel platforms.

Long temporal scales are required to capture chemical transformations in dynamic sampling events over large free-energy barriers. The energy landscape of realistic systems can be expected to be rough with a multitude of local minima and barriers. While intuition or experimental guidance can provide a useful bias, methods that can sample energy landscapes with minimal prior information are certainly desirable. Advanced algorithms adopted from other fields of mathematics and computing have recently shown some promise in achieving this goal. For example, the use of genetic algorithms to sample minimum-energy regions of configuration space (Hart et al. 2005; Abraham and Probert 2006; Oganov and Glass 2006) or of neural networks to perform unbiased fitting of complex potential energy surfaces (Behler 2008a, 2008b), combined with advanced sampling algorithms such as metadynamics (Laio and Parrinello 2002; Martoňák et al. 2003) or replica exchange dynamics (Sugita and Yuko 1999), are worthy of further investigation and optimization. One open question is whether accelerated dynamics approaches can be adapted to properly treat dynamics on excited states (see PRD titled, "Electronic Structure Theory and Simulation Methods for Nonadiabatic and Strongly Correlated Systems").

A novel variant of the event KMC method was recently demonstrated to dramatically extend the time scale accessible for simulations of microstructure evolution resulting from atomic diffusion (Oppelstrup et al. 2006). The new method of first-passage KMC is based on an exact factorization of the N-body diffusion-reaction problem and a time-dependent Green's function formalism (Oppelstrup et al. 2009). For a simple model of alpha-iron under electron irradiation, the method was shown to handle microstructure time evolutions spanning time scales from 10^{-11} seconds to hundreds of years (10^{10} seconds), without any compromise of simulation accuracy (Donev et al. 2010). The method is being currently extended to lattice gas models of multicomponent alloys where Green's functions, used for exact acceleration, are not available analytically. The first-passage KMC method and related methods under development achieve exceptional computational efficiency by focusing computational effort on the important infrequent events. The mentioned new methods are exact for an important class of material models, but further extension of the new event-driven simulation methods to more complicated (and realistic) material models requires trading some accuracy for efficiency. How to control error in such simulations is a topic of current research. Another difficult unresolved issue is whether it is possible to efficiently harness the immense computing power of platforms with massive multilevel parallelism for smart event-driven simulations like first-passage KMC. Recent general arguments suggest that it should be possible (Korniss et al. 2003), but specific development strategies remain unclear.

The basic idea of metadynamics is to activate a system to climb out of any potential well by means of a series of activation-relaxation steps. Originally developed for escaping from free-energy minima (Laio and Parrinello 2002), the method has been extended to probe slow dynamics in bulk systems undergoing structural relaxation. Results have been obtained for the temperature dependence of the shear viscosity of glassy liquids, where the viscosity varies by 15 orders of magnitude (Kushima et al. 2009). Preliminary results on thermal creep showing the stress relaxation of a model nanocrystal at constant

strain on the time scales of seconds and years suggest a broad range of rare-event or slow strain-rate applications (Lau et al. 2010). This method can benefit from access to large-scale computing resources.

Figure 12 shows the temporal variation of the von Mises stress in a model iron nanocrystal subjected to four constant initial strains (Lau et al. 2010). The curves illustrate the classical behavior of anelastic relaxation (Nowick and Berry 1972). An interesting feature is that the relaxation time is larger for the larger strain, a behavior that may be counterintuitive. However, this result may be rationalized by regarding the stress relaxation as a self-organized response, mediated by the microstructure, as opposed to a direct response to an imposed deformation. A larger imposed strain would mean a larger region of the system involved in the self-organized relaxation, which in turn takes longer time for the response to manifest. These results are obtained by a metadynamics method of sampling transition-state pathways (Laio and Parrinello 2002; Laio and Gervasio 2008) adapted for the study of thermal creep (Lau et al. 2010). The adaptation was first benchmarked in a study of viscosity of glasses (Kushima et al. 2009). Together, these studies constitute an alternative approach to atomistic simulation of slow dynamics where one is able to reach macroscopic time scales while retaining details of atomic configurations associated with the slow dynamics. In the case of thermal creep, the availability of atomic configurations associated with discrete relaxations reveals an essentially universal role of unit process for plastic deformation in small volumes. It remains to be seen whether such an approach can be extended to large-scale simulation in modeling chemomechanical problems such as those delineated in the PRD titled, “Mechanistic Understanding of Environmental Degradation, Coupling Chemistry and Stress.”

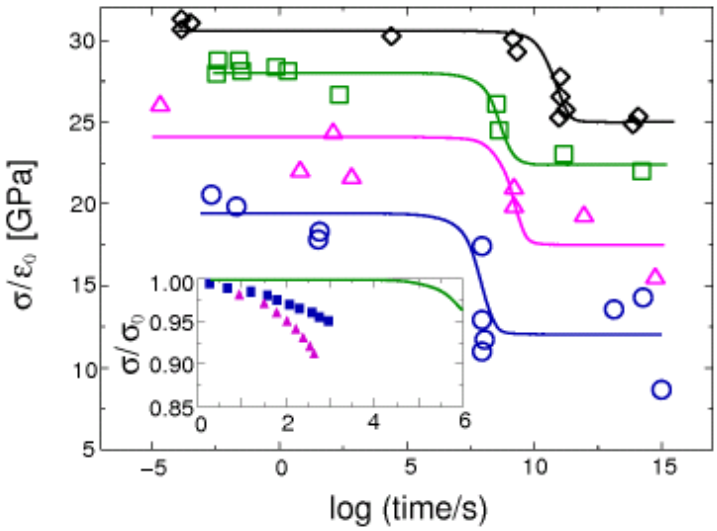

Figure 12. Temporal variation of the von Mises stress in a model iron nanocrystal subject to four constant initial strains, 0.005, 0.007, 0.010, 0.020 (circle, triangle, square, diamond, respectively) demonstrating anelastic relaxation (Lau et al. 2010).

Expected Computational and Scientific Outcomes

Advancing to the exascale thus offers a real opportunity for materials modeling. If researchers can fully harness 10^{18} flops, they can—in principle—perform atomistic simulations with a fast empirical potential, such as the embedded atom method in a domain where the product of length and time is 10^3 to 10^4 atom-seconds (e.g., a million-atom system for 10 milliseconds, or a 10,000-atom system for 1 second). This

opens the window to simulation of the abovementioned microstructural processes; these simulations can be used for direct investigation and understanding as well as for parameterizing and testing higher-level models to obtain the predictive power that is so vitally needed.

Specifically, researchers believe the following outcomes are achievable with the development of exascale computational resources and corresponding investment in algorithm development:

- atomistic simulation capability with full dynamical fidelity for reaching long times, even in the presence of low barriers
- improved understanding of processes underpinning materials microstructural evolution and response under various extreme conditions
- tools to parameterize and benchmark multiscale materials models
- algorithms/codes for massive multilevel parallelization of time and space.

Potential Impact on National Security

Accurate multiscale models for microstructural evolution on engineering time scales, including deformation and failure under extreme dynamic load and/or hostile environmental conditions are needed to meet nuclear energy and nuclear weapons certification challenges. As emphasized in this PRD, extending the time scale of calculations is even more challenging than extending length scales because of the intrinsically sequential character of temporal evolution. Thus, progress in advancing atomistic simulations on engineering time scales will allow more accurate predictions of material behavior under extreme conditions relevant to national security needs, which include the following:

- full microstructure simulations of dynamic material response to strain rates as high as 10^8
- full simulation of rapid solidification under pressure
- full simulation of single cascade annealing following radiation damage events.

In addition to these specific simulation needs, accurate multiscale models of microstructural evolution will provide robust benchmarks for semi-empirical approaches that extend to much longer spatial and especially temporal scales relevant to materials lifetime prediction and related challenges.

Predict Constitutive Representation in Multiphase Materials

Scientific and Computational Challenge

In this PRD, the scientific challenge is to provide an accurate representation of constitutive behavior of real materials under dynamic loading, including mechanisms of dissipation and damage in systems with realistic representations of materials heterogeneity.

The materials whose behavior scientists need to accurately describe are not perfect, homogeneous single crystals, but rather materials that have microstructure, defects, and interfaces. Further, these defects and interfaces evolve dynamically in extreme environments such as under compression and have

consequences (e.g., materials strength) at multiple temporal and spatial scales. Scientists must be able to predict the microscopic state of such materials as a function of time and from this information predict its constitutive response.

In the last decade, scientists have seen impressive breakthroughs in computational capabilities, and numerical simulations have provided exciting results at an unprecedented rate. Because of the spatial and temporal resolution that is now possible, interesting and unexpected phenomena have been observed at a variety of length scales. To make further progress in a predictive understanding of dynamic compression science, scientists must make progress on three parallel paths:

1. Develop experimental capabilities that provide time-resolved measurements at microscopic scales; i.e., researchers need to "look" inside the materials at length scales comparable to the mechanistically relevant length scales in numerical simulations.

2. Undertake numerical simulations that incorporate realistic dimensions and realistic microstructure, and are grounded in atomistically faithful, first-principles models.

3. Perform credible and quantitative comparisons between simulation output and experimental data to validate the robustness of the models.

A central challenge for NNSA is understanding materials behavior under extremes of strain and strain rate, up to and including materials failure. Researchers' limited understanding of the stochastic way in which polycrystalline metallic materials deform, damage, and fail prevents them from predicting the performance of these materials or preventing their failure. In high explosives, predicting the initiation and evolution of detonation and deflagration is similarly challenging. By understanding how these processes occur, researchers will be able to predict metallic materials responses in service and design new materials to prevent failure under specific conditions. These challenges are most acute in the presence of phase changes, under very large deformation, and when the scale of the inhomogeneity is comparable to the feature size of the component. It is also clear that materials inhomogeneities—due to, for example, alloying and variations in microstructure from processing changes—are key contributors to researchers' lack of predictive understanding. Advanced simulations enabled by petascale computations are beginning to provide relevant information on grain-scale inhomogeneity. Opportunities at the exascale, both in terms of the number of atoms that can be modeled and the time scales over which their evolution can be followed (see PRD, "Atomistic Simulation on Engineering Time Scales"), hold significant promise for qualitative advances, especially if these new calculations can be validated through experimental diagnostics with the same temporal and spatial resolution.

One of researchers' daunting challenges is that the heterogeneous microscopic structure of a material completely rebuilds itself during a phase transformation, and under extreme conditions researchers cannot complete the experiment to say what this microstructure is. They cannot pick the pieces and look inside. Rather, researchers must develop nondestructive in situ experimental techniques and the modeling tools to understand them. As the fidelity of the multiphysics simulations approaches the length scale of the heterogeneous microstate, the material constitutive behavior will vary significantly from zone to zone, stressing the need for a statistical representation that captures the local microstate and its resultant behavior. Furthermore, the microstate itself changes during dynamic loading and long-time aging.

The material with which scientists have come closest to meeting this challenge is highly defect-free crystalline silicon. Normally, a material is designed to meet an engineering need within a specified tolerance. In fact, it is the defects in materials that give rise to many of the constitutive behaviors (e.g., the effort required to deform a material) researchers must represent to the continuum multiphysics simulation codes. Vacancy and interstitial point defects are created during irradiation and diffuse through the material, resulting in changes in shape and other long-time deleterious effects. Dislocations are line defects that enable a material to permanently change shape when deformed and real materials are aggregates of small crystals containing inclusions, second-phase particles, and an inhomogeneous distribution of alloying elements. The interfaces between these crystal grains (grain boundaries) are sites for nucleation of new phases and the failure and fracture of materials.

The combination of accurate physical models with high-performance computing (HPC) has led to a new paradigm for the design, development, certification, production, and surveillance of weapons systems. Most of these models cull materials property data from a variety of sources to construct deterministic simulations of multiple phenomena spanning structural, thermal, fluid, and shock mechanics. The material properties are generally treated as averaged, scalar values—i.e., "the" yield strength of brass—but most engineering materials are processed in ways that produce variations in their internal structure. These variations can occur on the microstructure scale (i.e., local differences in internal structure), on the component scale (i.e., from one region to another on the same part), from part to part (i.e., due to batch and lot variations), and over time as the material ages. Furthermore, these inhomogeneities can affect both properties and performance, producing uncertainty therein. Thus, it is often necessary to treat the process-structure-properties performance spectrum statistically—and although this approach is usually neglected, it can be critical to a truly predictive capability for assessing performance and reliability.

When a component's dimensions are comparable to the microstructural features in the constituent materials (e.g., in some microelectro-mechanical systems), the variability of material properties can drastically affect basic design quantities. For applications with dimensions spanning only a few grains, the local orientation of the crystalline microstructure can dominate properties. Furthermore, localized critical events, such as ductile crack nucleation and growth, can be particularly susceptible to the randomness of a materials microstructure. In ductile failure, crack incubation, nucleation, and growth typically occur by void formation and coalescence. In this case, again the microstructure local to the material defect will dictate the peak load to crack initiation.

For ultra high-reliability applications, the uncertainty in properties is compounded by a lack of data describing the extremely rare events that are often critical to the design criteria. Physical testing alone is often unable to supply sufficient data. For example, if reliability requirements demand no more than a 1% probability of failure, then "proof" from physical testing would require at least a hundred tests (and likely many more). Obviously, with realistic reliability requirements, performing a very large number of physical tests is usually prohibitive. Ideally, physical testing would be supplemented with a very large number of accurate, digital realizations that can be simulated (relatively) quickly and efficiently.

To date, there is no robust, coherent approach to quantifying this uncertainty at the scale of materials microstructures, and then to using it in a predictive manner at the component length scale. To comprehensively quantify a component's performance and reliability, deterministic models of its materials physics must be combined with stochastic treatments of the materials microstructures to reflect the inherently statistical nature of the real materials properties. The inherent variability introduces additional requirements for existing HPC, partly because the need to sufficiently populate the statistics

could require computations that exceed existing requirements for "deterministic simulations" by several orders of magnitude. By exercising the United States' immense resources for high-performance, high-fidelity, deterministic simulations in a context that captures the inherent variability in both present and future engineering materials, researchers can establish a predictive, science-based capability that accurately captures not only the materials properties themselves, but also the statistics that underpin those properties. Only then can researchers hope to robustly address the sources of variability that produce uncertainty in the performance and reliability of the systems that are critical to national security.

Matching microstructure between the simulation cell and the experimental specimen is a daunting challenge in providing an accurate description of constitutive behavior of real materials under dynamic loading. The problem lies in choosing (in simulation) or characterizing (in experiment) the initial microstate of the material. Real engineering materials are processed far from equilibrium so that no global minimization procedure brings the sample to the "correct" macrostate. In most current computations of shock loading, pure single crystal samples are employed for which only one, trivial microstate is available (Luo et al. 2009). Even the inclusion of simple defects, such as twin boundaries, significantly alters the response (Srinivasan and Baskes 2007). This problem was already noted in comparing simulation with experiment in the simple case of thermal creep in a nanocrystalline iron (see Figure 12 [Lau et al. 2010]). Another particularly striking example of this computational difficulty concerns the so-called inverse Hall-Petch behavior that has been widely reported, both experimentally (Chokshi et al. 1989) and reproduced in molecular-dynamics simulations (Van Swygenhoven and Caro 1997; Schiøtz et al. 1998; Schiøtz and Jacobsen 2003). However, the experiments often suffer from difficulties with sample preparation, reproducibility of results, and control of grain size (Koch and Narayan 2000). In fact, experiments now show a diversity of behaviors at small grain sizes including "inverse" Hall-Petch, normal Hall-Petch, and a saturation of hardness that is independent of grain size. The molecular-dynamics simulations suffer similar difficulties as illustrated in Figure 13. In this work, nanocrystalline samples were annealed at 1100 K before high-strain rate deformation at 100 K. The figure plots the flow stress calculated by molecular dynamics as a function of the time at which the samples were annealed (Vo et al. 2008). Before annealing, the inverse Hall-Petch effect is clearly visible, and the flow stress increases with increasing grain size. After annealing and relaxation of the grain boundaries with no significant change in grain size, the inverse Hall-Petch behavior vanishes. In real engineering materials, the complexity is far greater because the materials are typically multiphase, multicomponent structures that contain dislocations, grain boundaries, and other defects—and all of these microstructural features are far from equilibrium. It will indeed be a grand scientific challenge to learn how to represent the initial microstate of such systems in making reliable predictions of materials response.

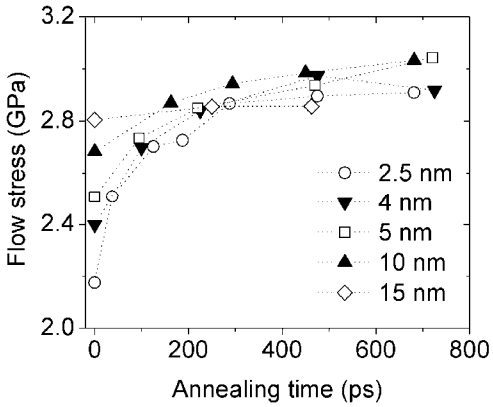

Figure 13. Flow stress as a function of grain size and annealing time at 1100K. Image courtesy of R. S. Averback (University of Illinois, Urbana-Champaign).

Summary of Research Direction

The research directions necessary to realize the needed statistical continuum representation of constitutive response in multiphase materials include the following:

- realistic representation of time-dependent material behavior at multiple length scales
- realistic representation of experimentally determined material heterogeneity at all scales
- methods that link in situ experiments and simulations at appropriate time and length scales
- mesoscale simulation models that capture the essential phenomena and bridge the continuum scale
- validation of multiscale models with simulations that overlap in time and space
- homogenization of direct numerical simulations using mesoscale models to create statistical continuum representations
- statistical mechanical and uncertainty quantification methodologies that faithfully represent inherent variability in materials properties and performance.

Expected Computational and Scientific Outcomes

Some of the expected scientific and computational outcomes enabled by exascale computing are as follows:

- simulations of material behavior with realistic microstructure—for example, the first molecular-dynamics simulations at large length scales (microns) compared to experimentally observed dislocation microstructures and at time scales (microseconds) relevant to dynamic experiments (e.g., shockwave loading)

- understanding the mechanisms of dissipation and damage during dynamic loading (e.g., entropy production)

- a detailed window into the kinetics of phase transformation, especially at conditions far from equilibrium

- calculation of in situ experimental observables and an understanding of microstructure development under extreme conditions

- validated mesoscale simulations that link the continuum scale.

Potential Impact on National Security

As previously noted, a central challenge for NNSA is understanding materials behavior under extremes of strain and strain rate, up to and including materials failure. Researchers' limited understanding of the stochastic way in which polycrystalline metallic materials—such as plutonium or uranium—deform, damage, and fail prevents them from predicting material performance or preventing failure. In high explosives, predicting the initiation and evolution of detonation and deflagration is similarly challenging. By understanding how these processes occur, scientists will be able to predict material responses in service and design new materials to prevent failure under specific conditions. Specific potential national security impacts include the following:

- multiphase, process-aware, path-aware constitutive models for use in multiphysics simulations

- the construction of inertial confinement fusion targets that create the optimal conditions for implosion—and more broadly, the predictive design of new materials for a wide array of desired functionalities and performance requirements

- the quantification of uncertainty due to inherent variation in material behavior for certification of the nation's nuclear stockpile, especially as it begins reducing the size of its arsenal.

Electronic Structure Theory and Simulation Methods for Nonadiabatic and Strongly Correlated Systems

Scientific and Computational Challenge

In this PRD, the scientific challenge is the predictive treatment of strong electron correlation and electronic excitation with quantum simulation codes for ambient and extreme conditions.

The behavior and phenomena that researchers need to compute are intrinsically quantum mechanical, and classical approaches consistently fail to provide accurate results. In the context of national security, examples where quantum mechanical treatments are essential include the electronic structure, equation of state, and constitutive response of F-electron systems such as uranium and plutonium; chemically reactive systems such as interfacial corrosion and initiation of high explosives; and strongly driven systems and ultrafast processes where nonadiabatic phenomena are central. In many of these cases, the computational challenge is not only faster calculations but also the capability to complete many calculations in parallel to explore a number of potential crystal structures, surface states, and electronic configurations as a means of accelerating progress.

Summary of Research Directions

Methods for treating correlated electron and nonadiabatic systems have been under development for decades, and it is tempting to conclude their treatment is an intractable problem. However, a combination of theoretical, algorithmic, and computational hardware improvements in recent years has made steady progress in combining "first principles" and "correlated electron" quantitatively. Continuing this progress by crossing the static threshold to correlated dynamics would be a breakthrough.

It might be assumed these intrinsically quantum mechanical problems are not amenable to solution with exascale computation resources because "scale" is generally associated with the number of atoms in a simulation, which are often not large for the classes of problems researchers are considering. However, in correlated-electron problems, "scale" has more to do with complexity than atom count—and thus the problems that could be addressed on an exascale platform are significant. Two specific examples are dynamic simulation of melting in plutonium and the properties of liquid plutonium under pressure using forces obtained from correlated electronic structure. Realizing these goals would provide information currently unavailable, yet important in determining the equation of state of plutonium, and provide a basis for understanding anomalous melting and strongly correlated liquids. Dynamic simulation of heavy metals with a normal correlation treatment is just possible with petascale capability; exascale capability would enable a predictive treatment of dynamics in actinides, providing the first demonstration of correlated electron dynamics.

Similar to the challenge of strongly correlated electron interactions is the simultaneous treatment of coupled electronic and ionic systems. Situations in which the dynamics of both nuclear and (hot) electrons are important, which are prevalent when one considers strongly driven systems or ultrafast processes, point toward the need for a theoretical and computational framework to treat nonequilibrium processes in the electronic and ionic systems simultaneously.

The Born-Oppenheimer approximation is the keystone to atomistic molecular-dynamics simulations. It is based on the assumption that the electronic wave function can readjust instantaneously on the time scale of nuclear motion. Electron and nuclear degrees of freedom are decoupled, and one can treat the former quantum-mechanically and the latter classically, opening the route to *ab initio* molecular-dynamics simulations. In the Born-Oppenheimer approximation, it is assumed that as nuclei move, the quantum-mechanical electronic subsystem continually remains in a state of equilibrium. Keeping this assumption in molecular-dynamics calculations is highly desirable, because it allows for larger and longer simulations.

There are situations, however, in which the Born-Oppenheimer approximation fails and the electronic and nuclear motion become coupled, giving rise to new phenomena. This happens, for example, when the nuclear velocities are large (in bond breaking during a chemical reaction, in highly energetic ionic collisions, etc.), when the dynamics are not governed only by the lowest energy electronic state (photochemistry, Zener tunneling, etc.), or when the exchange of energy between electrons and nuclei is important (Joule heating, electronic stopping power, etc.). In recent times, there have more and more attempts at incorporating quantum effects in the nuclear motion (e.g., tunneling or zero-point motion and electron-nuclear correlations), although these attempts have been confronted by the requirement of enormous computational effort.

The treatment of nonadiabatic effects in molecular-dynamics simulations has a long history. The simplest approach is the Ehrenfest approximation, in which the nuclei follow classical trajectories determined from

an average distribution of electrons. In spite of its limitations (failure of detailed balance), in recent years this method has become popular in the study of dynamics in systems with electronic excitations (Krasheninnikov et al. 2007) in combination with time-dependent density function theory. Probably the most widely applied method in chemistry is "surface-hopping" (Tully 1990) with its many variants. In this method, the forces on the nuclei are determined from single electronic potential energy surfaces, but hops between surfaces are allowed to include nonadiabatic effects.

A different approach is the electron nuclear dynamics theory (Deumens et al. 1994). The electron nuclear dynamics theory is a time-dependent nonadiabatic approach that employs the phase space view that emerges from the use of coherent states of the Ehrenfest representation. It has been successfully applied to a variety of ion-atom and ion-molecule reactive collisions. The correlated electron-ion dynamics method is a step beyond the Ehrenfest approximation, which aims at reinstating the electron-nuclear correlations, missing in Ehrenfest dynamics. The method relies on expansions of the quantum Liouville equation for the electron-ion system, with different formulations proposed in the limits of weak (Horsfield et al. 2005) and strong (Stella et al. 2007) electron-nuclear coupling. Although the correlated electron-ion dynamics method is still rather young and costly, it has emerged as a powerful tool for problems in which the transfer of energy between electrons and nuclei is crucial.

Due to the complications all of these nonadiabatic molecular-dynamics approaches introduce and the computational costs they require, their application to theoretical studies of materials under extremes remains challenging. The continuous increase in computational power and the improvements in the simulation methods and codes open new perspectives in the understanding of nonadiabatic phenomena in irradiated materials.

In addition to materials behavior in strongly driven situations, nonadiabatic processes also play a key role in chemically reactive environments where an accurate description of excited state phenomena is essential. Many important chemical transformations, particularly those involving concerted electron transfers as occur across interfaces, involve excited states in the reaction coordinate that is degenerate with the ground state. Such so-called "conical intersections" represent a difficulty in performing dynamics simulations because the system may oscillate resonantly between the ground and excited states. Current methods generally have difficulties dealing with excited states and dealing with level crossings at degeneracy. In addition, reactions that proceed via concerted multicenter, multielectron dynamics cannot be expressed by a Born-Oppenheimer surface. Such reactions are important, for example, in photosynthesis and as such are relevant for engineered photoelectrocatalytic fuel production systems. Proper treatment of these nonadiabatic dynamical processes in an efficient manner will likely require new theories, methods, and algorithms. In particular, efficient ways to simultaneously treat the dynamics of electronic and ionic degrees overcome the three-orders-of-magnitude difference in time scales are needed.

Expected Computational and Scientific Outcomes

Some of the expected computational and scientific outcomes enabled by exascale computation for intrinsically quantum mechanical systems include the following:

- Quantum simulation of the properties of materials, such as f-electron systems, whose interactions cannot be reliably predicted with classical interatomic potentials or with conventional electronic structure methodologies

- Accurate treatment of systems in which the lattice and electron degrees of freedom are strongly coupled and evolve nonadiabatically

- Electronic structure methods that faithfully describe excited state and dynamical processes that are central to chemical reactivity.

Potential Impact on National Security

Advances in nonadiabatic and strongly correlated electronic structure theory would enable predictive treatment of the following phenomena that are central to NNSA materials challenges:

- actinide electronic structure and equation of state under extreme conditions

- energetic materials and chemically reactive systems

- strongly driven systems/ultrafast processes.

More generally, the advances in atomically faithful quantum-mechanical simulation capabilities would directly affect the application areas described in other PRDs in this section such as interfacial chemomechanics and its consequences for materials compatibility, environmental degradation, and radiolytic corrosion; predictive constitutive representation of multiphase materials; and the robust coupling of advanced equation of state descriptions to strength and damage models.

CONCLUSIONS

The overarching grand challenge for materials research is to achieve a predictive, mechanistic understanding of real materials. In these circumstances, microstructure and interfaces matter, kinetics matter, and environments matter. Furthermore, the ability or inability to predict materials behaviors in these circumstances directly affects researchers' confidence or lack thereof in nuclear weapons performance. With the advent of exascale computing, the possibility now exists to achieve predictive capabilities to manipulate microstructure and interfaces at the grain scale to enable the design and development of extreme-environment-tolerant advanced materials. The exascale challenge, in essence, is direct access to materials functionality at the mesoscale frontier and prediction of mechanisms governing materials properties in micron-scale volumes over times of milliseconds. While framed in this panel report as specific to NNSA, these issues are a microcosm of broader materials challenges, especially in extremes such as materials needs for nuclear energy.

Table 2 lists the milestones for the work described in this section. Provided that the computational resources become available for research at the anticipated scales, the forefront research activities that will be conducted are provided as milestones.

Table 2. Computational Milestones for First-Principles Materials Simulations

Scale	Milestone
>1 Petaflop-year	• MD simulation of radiation-damage annealing, 10^6 atoms, 10^{-3} s • QMD two-phase simulation of high-pressure plutonium melt • DFT calculation of phonon spectra of strongly correlated electron intermetallic compounds (e.g., $CeCoIn_5$)
>20 Petaflop-years	• MD simulation of microstructural evolution • MD simulation of nucleation of phase transformations • DMFT calculation of relative structural stability of plutonium low-pressure solid phases
>100 Petaflop-years	• MD simulation of radiation-damage annealing, 10^5 to 10^6 atoms, 1 s • MD simulation of shock dynamics with realistic microstructure • DFT-MD simulation of a 1-keV radiation damage cascade • QMD simulation of plutonium low-pressure melt • SCP calculation of temperature-dependent phonon spectra of strongly correlated electron intermetallic compounds (e.g., $CeCoIn_5$)
>1 Exaflop-year	• Direct MD simulation of interacting cascades, 10^7 atoms, 1 s • MD simulation of shock plasticity localizing into macroscopic shear bands • MD simulation of nucleation and growth of a fission gas bubble in UO_2 • DFT-MD simulation of a 10-keV radiation damage cascade • DMFT-SCP calculation of temperature-dependent phonon spectra of low-pressure plutonium solid phases (hence the solid-phase diagram). • DMFT-SCP temperature-dependent phonon spectra of strongly correlated electron intermetallic compounds (e.g., $CeCoIn_5$) with accurate electron correlation • DMFT-QMD calculation of a low-pressure liquid isotherm.
DFT = density functional theory; DMFT = dynamical mean field theory; MD = molecular dynamics; QMD = quantum molecular dynamics; SCP = self-consistent phonon *MD simulations may be accelerated by a factor of 10-100 by expected developments in exponential accelerated MD methods.	

CHEMISTRY

Co-Leads: Robert Harrison, Oak Ridge National Laboratory
Steve Plimpton, Sandia National Laboratories

Panel Members: Matt Challacombe, Los Alamos National Laboratory; Bert De Jong, Pacific Northwest National Laboratory; Larry Fried, Lawrence Livermore National Laboratory; Giulia Galli, University of California, Davis; Edward Kober, Los Alamos National Laboratory; Felice Lightstone, Lawrence Livermore National Laboratory; Thomas Mattson, Sandia National Laboratories; Normand Modine, Sandia National Laboratories; Antonio Redondo, Los Alamos National Laboratory; Alejandro Strachan, Purdue University; Aidan Thompson, Sandia National Laboratories; Adri Van Duin, Pennsylvania State University; and Al Wagner, Argonne National Laboratory

CURRENT STATUS

Chemistry effects play an important role in the U.S. Department of Energy's (DOE) National Nuclear Security Administration (NNSA) general interest in material properties in extreme environments. Chemistry, in the context of this panel report, is the breaking and forming of bonds between atoms. In computational models, this occurs at the quantum level, where electrons are included explicitly in density functional theory (DFT) or quantum-chemistry methods; at the atomic level in molecular dynamics and Monte Carlo methods that use interatomic potentials (IP), which include bond-order effects; and in mesoscale and even continuum models that include chemical reactions in their Hamiltonians or via rate equations.

Chemical reactions in materials commonly occur because inhomogeneities (e.g., a defect, a stress or temperature gradient, or an interface) are present or were created by a previous reaction. Consequently, an accurate model must include an adequate sampling of such features and the conditions that create them. This greatly expands the spatial extent and number of configurations that must be considered for simulations to be predictive, which—in turn—drives the need for large-scale computing.

Chemistry panel members identified and described outstanding scientific challenges mediated by chemistry effects that meet two additional criteria: 1) these challenges are within NNSA's charter or affect national security more generally; and 2) progress in the study of these challenges and eventual resolution could be significantly affected by the availability of next-generation extreme-scale computers.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

The five basic science challenges (BSCs) are listed below. The first three BSCs are of critical interest to DOE and NNSA; the last two are of broader national security interest.

1. Understand the response of target materials in inertial confinement fusion devices such as the National Ignition Facility (NIF) and the Z-pinch machine

2. Assess the safety, control, and performance of high explosives used in nuclear weapons

3. Predict the aging of weapon components

4. Design molecular sensors using first-principles modeling

5. Predict how actinides and their associated compounds will behave in environmental settings.

For each BSC, the following issues are highlighted in subsequent sections of the report:

- motivation for NNSA or national security
- current state-of-the-art technology
- science challenges
- need for exascale computing
- priority research directions (PRDs) to pursue.

Based on these science drivers, the following four crosscutting PRDs were identified. The expectation for these PRDs is that significant near-term progress can be achieved with the availability of computing resources at the extreme scale, and thus have a significant impact on the BSCs:

1. Advance the "speed and accuracy of quantum chemistry modeling" with the goal of 1 kcal/mol accuracy for ground states, transition states, and excited states all in condensed-phase environments. This will require development of new algorithms that perform well on extreme-scale platforms.

2. Enable "predictive actinide chemistry" for condensed-phase environments, including defects in solids and counter-ions in solution. This will require more accurate DFT functionals and equation-of-state calculations for f-electron materials.

3. Enable "reactive molecular dynamics with quantified accuracy," so dynamics can be modeled at the micron/microsecond scale with a chemical accuracy of 1 kcal/mol. This will require automated methods of generating reactive force fields for a broad range of materials by incorporating quantum-calculated information and training reactive models appropriately.

4. Bridge length and time scales to enable "seamless multiscale chemical reactivity" modeling for reactive systems. This will enable direct verification of the coupling algorithms used between methods and create a predictive modeling capability for high explosives. In environmental settings, it could enable predictive modeling of hydrodynamic flow under reactive conditions, a situation where it is extremely difficult to decouple the length and time scales of chemistry and physical properties or response.

Table 3 emphasizes the crosscutting nature of the four PRDs; each is relevant to several of the five BSCs.

Table 3. Crosscutting Issues Related to Priority Research Directions and Basic Scientific Challenges

	BSC 1	BSC 2	BSC 3	BSC 4	BSC 5
PRD 1	√	√	√	√	√
PRD 2	√	--	√	--	√
PRD 3	√	√	√	√	√
PRD 4	--	√	√	√	√

The individual BSC and PRD discussions detail specific needs for extreme-scale computing for chemistry-driven science problems. The following plot in Figure 14 illustrates the first of two recurring themes:

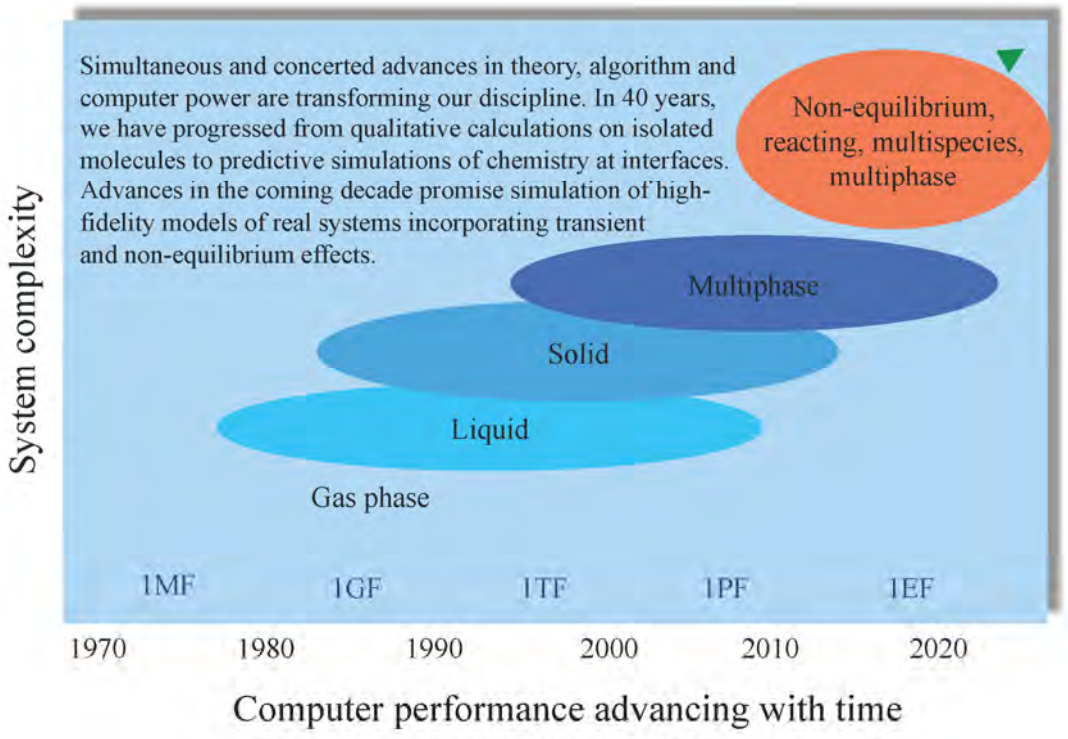

Figure 14. Scope of quantum chemistry modeling as a function of time and compute power. Image courtesy of Robert Harrison (Oak Ridge National Laboratory).

In a broad sense, the level of detail that can be accurately modeled in quantum chemistry calculations has progressed over the last three decades from gas-phase interactions between two small molecules to near-neighbor interactions of small molecules with solvent in the liquid phase to crystalline solids with large unit cells. This has been enabled not only by algorithmic advances but also by the availability of supercomputers with peak speeds increasing from gigaflop to teraflop and the current petaflop range. This level of computing is necessary because of the following:

- Model sizes have increased from a handful of atoms (gas-phase) to dozens of atoms (liquid) to many hundreds or thousands of atoms (solid).

- The computational cost of quantum methods typically does not achieve linear scaling in N, the number of electrons in the system, but is often closer to N^3.

The next frontier is interfacial or multiphase chemistry, where a predictive capability requires that quantum effects at solid-solid, liquid-solid, or gas-solid interfaces be modeled accurately. An additional challenge is to enable nonequilibrium processes, where rapid changes are occurring, to be modeled accurately. Many of the problems outlined in the BSCs are interfacial and nonequilibrium in nature. Examples include the mixing of target materials in a fusion implosion (BSC 1); heterogeneous high-explosive material composition at the nanoscale and microscale (BSC 2); changes at surfaces that mediate

aging effects (BSC 3); adsorption of molecules to a detector surface in a sensor (BSC 4); and adsorption onto surfaces of a porous media during groundwater migration of radioactive contaminants (BSC 5).

Because interfaces break symmetry in a computational model, such problems will require thousands or even tens of thousands of atoms to be modeled at a quantum level of detail with a new level of available computing power. Many of the models highlighted here will be possible to simulate for the first time on exascale platforms.

The second recurring theme is the development of new chemistry-driven models because the interfacial scale requires advances not only in computing power, but also in fundamental algorithms. This is needed to capture new physics (e.g., f-electron chemistry) and to couple across multiple length and time scales. Before their use in a predictive fashion, new algorithms require testing and validation. This testing and validation phase for the scale of NNSA and national-security science problems as discussed in this report will also require extreme-scale computing. Three aspects of this effort are as follows:

- Bond breaking and formation is, by its nature, "rare-event modeling." When these events trigger large responses, as in the case of detonation initiation (BSC 2) or aging (BSC 3), a large ensemble of calculations must be run to calculate statistically meaningful results if predictive accuracy is desired (such as to quantify safety margins on a material). This issue is compounded if the system itself is large, and its response can depend on stochastic variability in its initial state.
- New interaction potentials, models, and methods for extreme environments (BSC 1), reactive high explosives (BSC 3), and actinide chemistry (BSC 5) cannot be "validated against experiment" unless simulations can be run at large enough scale. For chemistry effects, this need is amplified because the models are more expensive to compute than nonreactive materials science models, and because tracking chemistry requires smaller time steps. Thus, the availability of faster computers will enable the ability to develop better models.
- When coupling across length and time scales with different methods (quantum to atomistic to mesoscale), the best way to validate the technique is to perform calculations with each method that overlaps; i.e., a large computation with the fine-grain method and a small computation with the coarse-grained method on the same system (the "seamless" approach of PRD, "Seamless Multiscale Chemical Reactivity"). For some of the science problems discussed in this panel report, it is anticipated such overlapping calculations also will be possible for the first time on extreme-scale platforms.

In summary, this report highlights two uses of extreme-scale computing for chemistry-driven models. The first is for selected heroic-scale computations to tackle the basic science challenges outlined. The second is to enable development of methods and models with a predictive level of accuracy—a task that still is an outstanding challenge for interfacial chemistry. The latter task will take advantage of an exascale platform to perform sufficiently large numbers of (mostly) independent calculations to parameterize and validate new techniques. This may not be the preferred use mode of such a machine in a programmatic sense, but it is the way new computational science is conducted.

The following sections outline the five chemistry BSCs and four associated PRDs.

Diamonds Are Not Forever

Carbon is found in two different structures under normal conditions: graphite and diamond. Although both are made of carbon atoms, the two different atomic structures result in vastly different properties of the materials: graphite is black and soft, diamond is transparent and hard, in fact one of the hardest materials known.

When faced with the challenge of designing a fusion capsule of diamond for the National Ignition Campaign (NIC) or deciphering the structure of a giant planet like Neptune, it is vital to know the behavior of carbon under conditions where the pressure is 10 million atmospheres (1000 GPa) and the temperature is several thousand degrees.

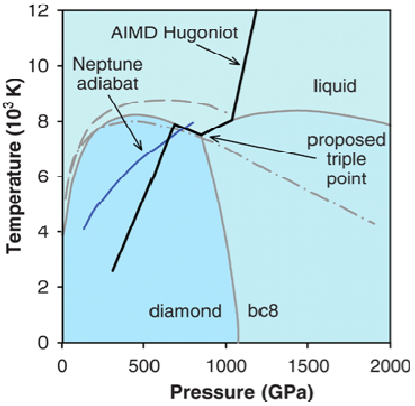

Phase diagram of carbon (gray solid lines are phase boundaries) showing the triple point of diamond, bc8, and liquid. The shock trajectory (Hugoniot) is the full black line (Knudson et al 2008).

Advanced quantum chemistry simulations such as those described in this BSC can predict which structures should exist. For carbon, the BC8 structure is favored over diamond at high pressure. BC8 is a tetragonal arrangement that packs atoms more efficiently than the diamond structure. In addition, the simulations showed that when subjecting diamond to shocks of increasing strength, the diamond begins to melt at 7900 K and 680 GPa and completes melting in a complicated way. A nearly 400 GPa solid-liquid coexistence region, starting around 680 GPa and completing around 1040 GPa, encompasses a diamond-liquid-BC8 triple point at 850 GPa. Such a complex behavior has profound implications for the behavior and design of a fusion capsule.

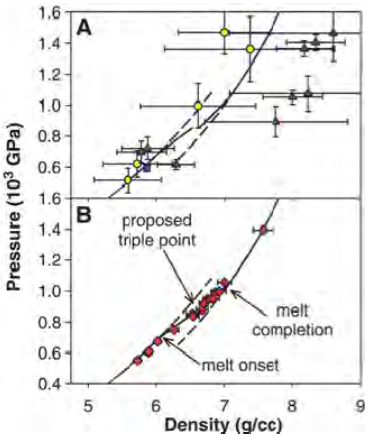

Shock pressure of diamond (Knudson et al 2008). Black full line is prediction from DFT/AIMD. Symbols are data using laser drivers (A) and magnetic flyer drive (B).

When the outcome of a fusion experiment lies in the balance, it is imperative to validate theoretical equations of state with accurate experiments. The second figure shows how well these simulations predicted the complex shock-melting of diamond (Knudson et al. 2008).

Integration of experiments and theory with high fidelity is the basis for understanding chemistry and physics under extreme conditions.

Images courtesy of Marcus Knudson (Sandia National Laboratories).

Material Response in High Energy-Density Regimes: Chemistry at the Extreme

Motivation for National Nuclear Security Administration

A thorough understanding of a wide range of materials subject to extreme conditions is critically important for many of NNSA's mission areas and a main topic of research at several major facilities. Experimental and theoretical efforts are tightly integrated within the inertial confinement fusion (ICF) programs at Lawrence Livermore National Laboratory's NIF and Sandia National Laboratories' Z-pinch machine, where extremely high-energy densities are induced as a small target implodes. Proton radiography (pRad) at the Los Alamos Neutron Science Center (LANSCE) is used to study the strength and spall of materials and properties of high explosives. The Dual Axis Radiographic Hydrodynamic

Test Facility (DARHT) at Los Alamos National Laboratory is used to analyze the dynamics of materials. Although these facilities are different in character, they share pressing mission needs to predict, design, optimize, and quantify uncertainties in materials behavior under extreme conditions.

A reason for the intense interest is that matter notably behaves differently under extreme conditions in density, temperature, and pressure. For example, seemingly basic elements and compounds—such as carbon, lithium, and water—display dramatic changes when exceeding normal conditions. With increasing pressure, water turns superionic (Cavazzoni et al. 1999; Goldman et al. 2005), a phase where the oxygen atoms are frozen into a lattice, while the hydrogen atoms move freely. With increasing temperature, water then turns into a metallic fluid (Mattsson and Desjarlais 2006), even taking part as a catalyst in chemical reactions under extreme conditions (Wu et al. 2009) as indicated by Figure 15. With increasing density, lithium displays an array of different solid phases with striking changes in melting temperature. When shocked, carbon in the diamond phase melts at a pressure of 9 million atmospheres as it simultaneously transitions to the solid BC8 phase (Knudson et al. 2008), a behavior that will affect the design of fusion capsules described in the sidebar, "Diamonds Are Not Forever."

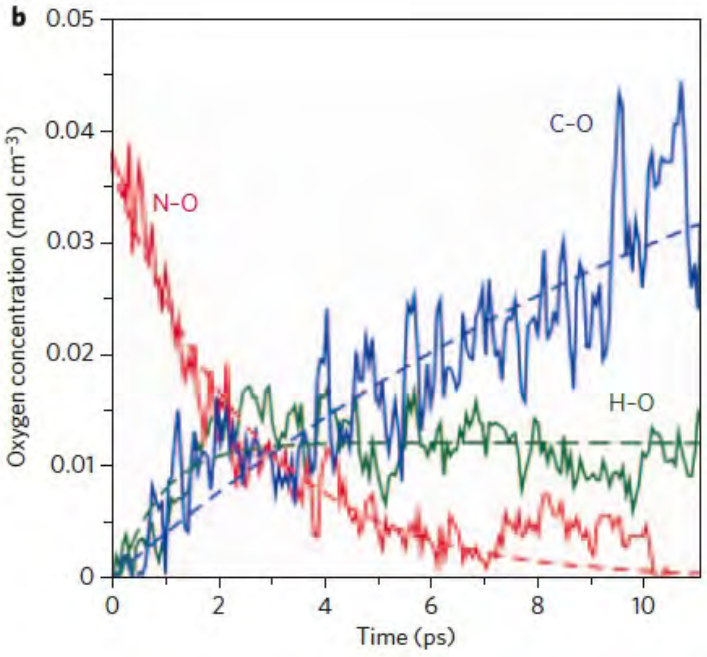

Figure 15. Time-dependence of the concentration of oxygen bonded with hydrogen (green), nitrogen (red), and carbon (blue) in a four-element system at 3000 K. Because C-O and N-O bonds increase and decrease with time, while the H-O bond count remains roughly constant, this suggests oxygen is flowing from nitrogen to carbon via catalytic reactions mediated by OH. Image courtesy of Christine Wu (Lawrence Livermore National Laboratory) and *Nature Chemistry*, 1, 57-62 (2009) (Wu et al. 2009).

These dramatic property changes make the transition from a largely descriptive treatment to a predictive capability a major step, which will require significant new advances in computational chemistry and computing capabilities. The nuclear test ban treaty and renewed commitment to a world free of nuclear weapons—"the path to zero" as envisioned by the current government administration—adds importance and urgency to this topic. The significant interest in nuclear energy as a path to increased energy security is also a major motivation for extensive work in the area of chemistry under extreme conditions.

Current State-of-the-Art Technology

Although the importance of chemistry is widely recognized, fundamental understanding of chemical processes requires application of quantum mechanics and quantum chemistry. Contemporary chemistry is fundamentally dependent on high-performance computing and the area of high-energy density science is no exception. The combined use of supercomputers and increasingly accurate theoretical models has yielded significant progress. In fact, chemistry at the extreme leads to challenges not present in gas-phase chemistry. Chemical reactions occur in a condensed, disordered, and fluctuating medium. Therefore, realistic modeling requires dynamic simulations of hundreds of atoms for several picoseconds. DFT-based molecular dynamics is the current method of choice. Traditional density- and density-gradient-based exchange-correlation functionals have the computational efficiency required to treat many hundreds of atoms for multiple applications with acceptable accuracy. This approach has resulted in breakthroughs in understanding the chemistry and physics in extreme environments.

Science Challenges

Chemistry and physics under high energy-density conditions has been called the X-games of contemporary physics (National Research Council 2003) because of the extreme challenges facing efforts in this area. Pressures at 10 million atmospheres (TPa) enable new phases of matter such as superionic water, glassy and eventually metallic hydrogen, and phase separation between otherwise soluble elements, including hydrogen and helium. Temperatures at thousands to tens of thousands of Kelvin create a strongly fluctuating environment where temperature not only affects chemical bonding, but reactions with rather high-reaction barriers will frequently occur. The main challenges in making the step from a largely descriptive treatment to a predictive capability fall into three categories, as described in the following sections.

Improved Fidelity of Electronic Structure Calculations

Although DFT is a theory with a 40-year history, it is by no means finalized. The exchange-correlation functionals in use today are largely split between hybrid functionals (B3LYP) in chemistry and gradient functionals (LDA, PW91, PBE, AM05, and WC) in solid-state physics (Mattsson et al. 2008). Exchange-correlation functionals have a long way to go to be able to treat mixed molecular and condensed systems, including actinides and transitions metal oxides, with chemical accuracy (1 kcal/mol). Chemistry under extreme conditions is an area where the two regions of applicability collide, and an improved way to calculate the quantum exchange-correlation effects is highly sought. Hence, fundamental work in DFT would leverage the power of new generations of computers and result in qualitatively more significant progress than can be reached if development work is limited to porting existing functionals to extreme computing platforms. Obtaining the wrong result faster is not the objective of the new computing platforms.

Treat Larger Systems with the Best Methods

Currently, computational resources directly limit the size (number of atoms) that can be included in simulations. For pure elements such as hydrogen, helium, lithium, beryllium, or carbon, available computers and methods have proven to be sufficient to allow for significant breakthroughs in understanding their behavior under extreme conditions. However, for many applications, the behavior of mixtures must be taken to the same level of predictive modeling. Understanding not only binary, but

ternary and higher mixtures, is essential for interpretation of experimental data. Treating low- to trace-concentration mixtures is an example of another demanding application. High-fidelity simulations of the chemistry of light elements, transitions metals, and actinides are imperative for understanding the dynamics in these systems. The complexity can be advanced further for studying phase transitions, including chemical reactions. Noncongruent transitions occur when during melting or vaporization, the chemical composition is different in the two phases. Such conditions require much larger computational systems to be investigated with high fidelity.

Quantify Uncertainties in Materials Behavior

The mission needs of predicting, designing, and optimizing materials behavior under extreme conditions require a systematic approach to uncertainty quantification (UQ). This effort will require a significant increase in computational requirements; e.g., for sensitivity studies.

Need for Exascale Computing

Already, first-principles simulations in high-energy density conditions are highly demanding at present goals. The field is currently capability limited. Modeling current system sizes using current methods with higher accuracy—e.g., hybrid or exact exchange/optimized effective potential methods (EXX/OEP)—will conservatively add a factor of 100 to 1000 to current computational demands. To improve the treatment of mixtures and phase transitions, a factor of 10 in system size resulting in a factor of 1000 in computational demand, would make a significant difference. Larger systems also require longer simulation times, which together with an increased emphasis on UQ, will add another factor of 100.

The area of high energy-density science is at a stage where dramatic improvement in fidelity with exascale computing is expected. The time to solve these problems is now. Experiments are costly (lasers at NIF and Laboratory for Laser Energetics; Z-pinch at Sandia National Laboratories) and achieve a limited set of conditions via a limited number of experiments. Simulations have a broader range of applicability. In addition, a clear interplay between theory/simulations and experiments leverages experimental outcomes, maximizing the return of the significant investments already made in experimental facilities.

Research Directions to Pursue

The science challenges described have direct consequences, which are important research directions to pursue. To improve the fidelity of electronic structure calculations, a main challenge lies in the treatment of actinides and transition metals. With all-electron (relativistic) molecular dynamics, it is possible to routinely treat heavy elements with the same accuracy as first-row elements. There are several technical reasons for the difficulties encountered in modeling actinides: the large number of electrons, the mixed core/valence character of f-electrons, relativistic effects, and shortcomings of existing density-functionals. To move the state of the art for actinide physics and chemistry to the same level as traditional first-row chemistry, close to all-electron molecular dynamics in combination with new functionals (traditional gradient-corrected or orbital-dependent/EXX/OEP) will be needed. If dramatically increased computational resources are coupled with new theory, the physics and chemistry of actinides is an area likely to reach a tipping point.

Efforts should be made to establish new and improved methods for first-principles calculations of transport properties, including viscosity, thermal conductivity, electronic conductivity, and opacity under extreme conditions. Predictive methods that seamlessly work for wide-gap insulators, semiconductors, and metals also need to be developed. Nonequilibrium methods (e.g., time-dependent DFT) are promising but require significant research before they could be widely used. Intrinsically, these investigations are computationally demanding—in particular for materials close to phase transitions.

The simulation of thermo-physical properties (such as phase transitions and separation) for low-concentration alloys and compounds set different, but equally significant, directions for research. Renewed efforts for efficient calculations of electronic structures for these large systems are of utmost importance. Most—if not all—DFT codes in production use today are based on algorithms that scale as N^3 or worse with the number of electrons. Therefore, even massive increases in computational power yield limited increases in system size. Methods that scale linearly in system size and can be executed on future heterogeneous parallel platforms would have a significant impact. However, it is imperative these methods allow for systems to be metallic and for treatment of thermal equilibrium. These two requirements are of particular importance to chemistry under extreme conditions because under high pressure and temperature, even water and hydrogen become metallic.

Control of Safety and Performance for High Explosives

Motivation for National Nuclear Security Administration

High explosives are both a critical driver component in the current nuclear and conventional stockpile and a major safety consideration. Current confidence in these materials and their handling is bound by extensive experimental measurement and empirical correlations. Any adjustments of the current formulations (driven by manufacturing considerations or occurring through aging processes) must undergo significant testing due to insufficient confidence in models for these systems. The development of new formulations is an expensive and arduous process, particularly for safety assessments during advances in the formulation processing. The assessment of unusual formulations used by terrorist organizations presents additional challenges because of large variability associated with their processing. Consequently, there is a significant common need from NNSA, the U.S. Department of Defense (DoD), and the U.S. Department of Homeland Security (DHS) agencies to dramatically improve current understanding of high explosive formulations, especially regarding safety and performance. Exascale computing can be used to address several current problems in this arena.

Current State-of-the-Art Technology

Current performance and safety models are largely empirically based and heavily calibrated to experimental tests. Although high explosives are heterogeneous composite materials, they are typically treated as homogeneous materials in continuum codes with homogeneous temperature and stress distributions (Menikoff and Shaw 2008). Because initiation and safety phenomena are controlled by localized fluctuations, these models are not directly tied to limited fundamental understanding and cannot be relied upon outside of their calibration regime. For safety considerations in particular, there are a large number of accidental initiation scenarios that must be considered (manufacturing and processing, damage from dropping or mechanical insult, nearby flame or explosion, and/or electrostatic discharge). With only limited understanding of the inherent fluctuations in these complex scenarios, experiment alone cannot confidently contain the risk. Development of improved models for these conditions is one of the major

challenges for the current NNSA Advanced Simulation and Computing (ASC) Program and has attracted growing interest from the DoD community.

These issues are primarily addressed at the mesoscale level, where the heterogeneous nature of the materials is explicitly represented. These models typically include the crystalline explosive base (~100 micron crystals, typically quite brittle), a polymeric binder (may be present in submicron thicknesses with highly temperature-dependent viscoelastic properties), and voids of various sizes (both internal to the crystals and between the crystals and binder). Analytic evaluations indicate voids are dominant focusing agents that can instigate ignition, and resulting "hot spots" of dimension ~0.1 micron are required for self-sustaining reactions. Experimental characterization of these materials over this range of length scales is only now becoming accessible, and characterization of the dynamic damage evolution processes is a current grand challenge (see Berghout et al. 2006; McGrane et al. 2009; and Smilowitz et al. 2008). Many damage evolution models are based on the concept of fracture of the crystals by an initial shock wave followed by frictional heating caused by sliding between those surfaces. Accurate modeling of the fracture, frictional heat generation, thermal conductivity, phase transitions (melting), chemical reactions, and gas generation should be incorporated within a realistic geometry for a successful evaluation. This would require especially large-scale, three-dimensional simulations with an accurate description of phenomena at interfaces. Most current work is associated with two-dimensional simulations on model geometries as a means to prioritize the various aspects for further study. Where limited data are available, accurate material properties over a broad range of temperature and pressures are also required for this effort because of the inherent instability of the materials.

At this level of simulation, limitations still exist due to the attempted description of interfaces, usually as well-defined, smooth surfaces. For a reacting high explosive, there is a pronounced blurring between a well-ordered nonreacting crystal, through a likely disordered reacting condensed phase, to a primarily product-based gas phase where reaction is likely to complete. Atomistic simulations can—in principle—provide an explicit representation of these various phenomena and link them to chemical reactivity (Cawkwell et al. 2008; Zhang et al. 2009). An example of such a "quasi" two-dimensional calculation with 17 million atoms is shown in Figure 16. However, these models cannot yet reach the realistic length scales required to characterize the microstructure and dynamic response. For example, electronic structure calculations on ensembles of hundreds of atoms are currently used to characterize the equation of state of both the reactants and primary products. However, clusters of solid products (soot and/or metal oxides) are likely much larger than this and cannot be included in the analysis. Yet, they are a significant lever arm in understanding the equation of state. Similarly, reactive molecular dynamics simulations now can be performed on many millions of atoms. This provides a basic understanding of the chemical reactions but is not sufficient to understand propagation or coupling to gradients in stress or composition. Finally, nonreactive molecular dynamics models of 100 million atoms (~0.001 micron^3) are being used to understand deformation processes and model the collapse of a single nanoscale void, but they are insufficient to understand the interactions between multiple voids at the micron scale.

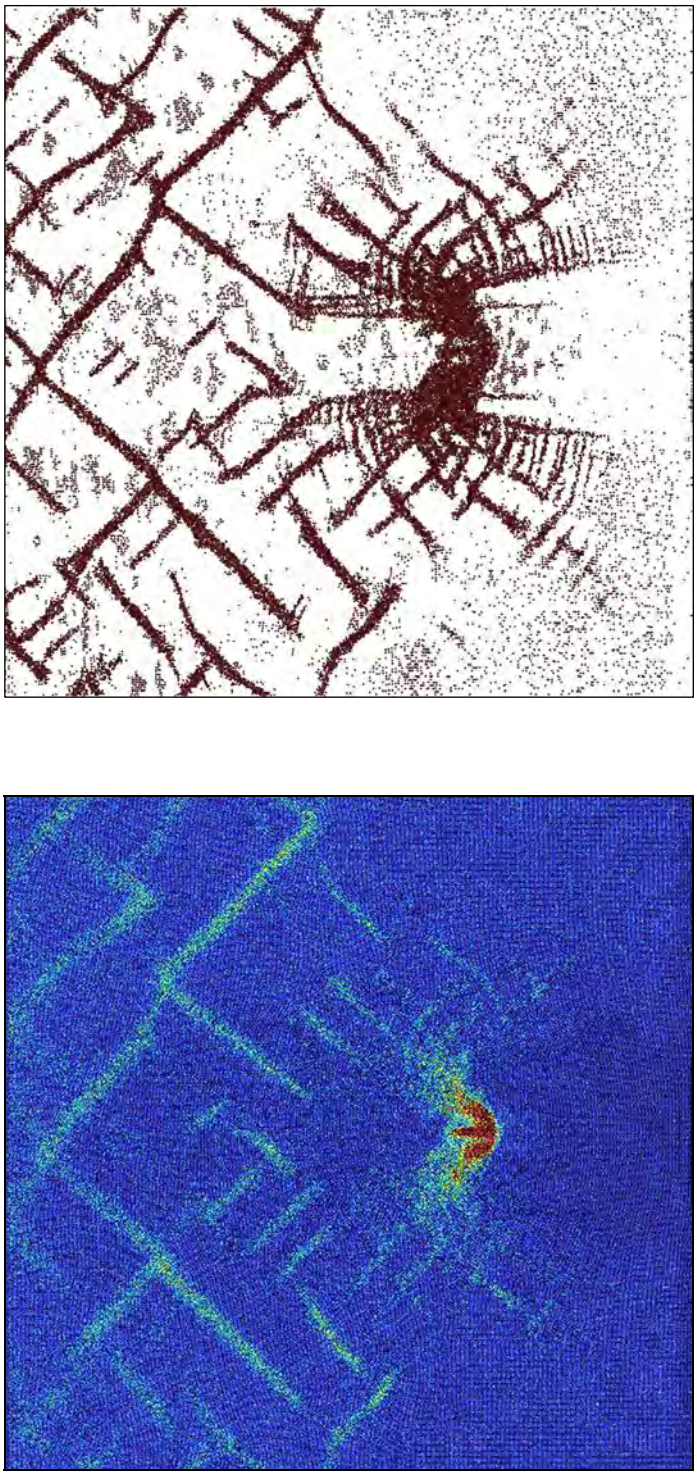

Figure 16. Top image shows response of a quasi two-dimensional sample (300 x 2 x 250 nanometers) to a plate impact from the left. A 50-nm-diameter void initially exists slightly to the right of the center point. The top image shows shear banding and conformational changes in the bulk material before interacting with the void. The bottom image shows the strong rise the temperature field as a result of the void collapse, which might lead to detonation initiation. Images courtesy of Marc Cawkwell (Los Alamos National Laboratory) and Thomas D. Sewell (University of Missouri-Columbia).

Science Challenges

The core challenge presented here is to couple chemical reactions and hydrodynamics in an inhomogeneous system. The length scales of these phenomena overlap, so they cannot be decoupled but must be treated as an integrated whole. Further, this happens in a condensed phase system at high temperatures and pressures where there is limited experimental capability to probe the system or maintain steady-state conditions. Consequently, calculations must supply much of the information that cannot be obtained otherwise, which requires well-validated models that can be linked across length and time scales. Atomistic simulations can provide useful guidance on these issues and are certainly the only way to obtain an explicit representation of chemical reactivity, but they cannot yet reach the realistic length scales required to characterize the microstructure. Mesoscale (polycrystal) simulations are needed to include the microstructure effects, but they must be performed at high resolution to capture the prominent gradients near reacting interfaces. The accurate coupling between the chemistry and microstructure of these materials is critical for the development of predictive models.

Need for Exascale Computing

As discussed, various types of atomistic simulations currently are providing insights into various issues. By extending the accessible length and time scales, exascale computing will enable integrated solutions that build on the successes of atomistic simulations.

This scale of computing is also required for reducing and bounding the uncertainties and fluctuations of the simulations. This can be applied to three distinct classes of simulation outlined above: 1) electronic structure, 2) reactive molecular dynamics, and 3) nonreactive molecular dynamics.

Electronic structure calculations provide the base model for these systems and the most direct tie to calibrating experimental data. As previously mentioned, calculations on systems containing in excess of thousands of atoms are needed for the calculations of equation of state because of the presence of solid product clusters. In addition, the results from thousands of randomly selected configurations are needed to supply statistical significance to these results. In addition, electronic structure calculations supply much of the base information on chemical reactions. These are typically performed on gas-phase uni-molecular or bi-molecular configurations, while condensed-phase results at the appropriate density would be required for a more complete understanding. This will require an enormous number of calculations to establish a basis set for calibration.

Current reactive molecular dynamics simulations are limited to nanosecond time scales for millions of atoms using classical IP. The reaction zones for high-order detonation processes (the faster possible reaction rates) are estimated to be 10–100 microns in linear extent and 1–10 nanoseconds in temporal evolution. In addition, there are strong pressure and temperature gradients in the profiles as well. As such, analysis of a static box supplies insufficient information. The location of the sonic plane (where the local sound speed equals the propagation velocity) in this dynamic structure controls the detonation propagation characteristics, which depends on confinement and initial temperature. In addition, the interactions of this wave with chemical fluctuations, interfaces, and voids are unknown but suspected to play significant roles in refined models. Such simulations would require ~10 micron^3 of material or ~10^{12} atoms observed for time periods in excess of 10 nanoseconds. Initiation and deflagration processes extend out several orders of magnitude larger in both length and time scales.

Because it is computationally less intense, nonreactive molecular dynamics can extend these scales out an order of magnitude or two in both space (>100 micron^3) and time (>1 microsecond). Simulations at these scales can resolve the mechanical effects for multiple voids, stress waves propagating across binder layers, and crack propagation at moderate strain-rates. This information will resolve many issues that are inadequately addressed in current mesoscale models because of limited ability to characterize interfacial phenomena. These results will be needed to provide substantial closure relationships for those models.

In addition, high-explosive component aging in nuclear weapons is a growing concern. This essentially results in a slightly different formulation being present compared to what was initially installed. There is a parallel need for the development of new high-explosive formulations with specific properties that can serve as replacement materials in many current munitions designs. The development of entirely new formulations is also desirable but entails enormous overhead related to current methods of scale-up and safety analysis. Thus, the availability of exascale computing to reach the necessary length and time scales is a timely development.

Research Directions to Pursue

The classical molecular dynamics and mesoscale modeling simulations are adapted to parallel platforms with near linear scaling, and they (nominally) can achieve the extended length and time scales previously discussed. For molecular dynamics, some issues associated with the amount of local memory needed for complex potentials and communication issues required for long-range electrostatic interactions need to be addressed. For mesoscale models, higher-order methods that can accurately describe systems with strong gradients and more complex materials models that retain accurate treatment and knowledge of fluctuations (especially temperature) need to be developed. For the electronic structure calculations, better scaling algorithms (approaching linear scaling for both the number of atoms and the number of processors) are needed that will dramatically reduce time to solution for these large systems. A robust methodology to validate molecular dynamics potentials that can reproduce the electronic structure calculations (and experimental data) to some specified level of accuracy need to be developed. In addition, algorithms for identifying and extracting results from these massive simulations and performing instantaneous analyses are needed to avoid the storage and subsequent analysis of the data. This is particularly acute for the molecular dynamics simulations where time averages over localized regions of space are needed to extract thermophysical properties, such as temperature, stress state, and reaction rate. Isolated snapshots of the simulations cannot provide these data, and saving all data for post-analysis is impractical, if not impossible.

Predictive Chemistry of Aging

Motivation for National Nuclear Security Administration

As the nuclear weapon stockpile continues to reduce in pursuit of President Obama's goal of "a world without nuclear weapons," maintaining the nation's nuclear deterrence imposes more stringent requirements to confirm each remaining weapon performs as expected. Weapon performance in actual use is determined by the properties of weapons components at the time of application, which are the result of a process of evolution from the initial (as manufactured) structure and properties to a new, aged structure with changed properties that may or may not be consistent with the expected performance of the system. Therefore, the ability to predict the properties of components and their degradation over time is an essential requirement for maintaining nuclear deterrence. Furthermore, in an environment that

substantially curtails the ability to replace or refurbish aging weapons systems, the time scale over which DOE must understand aging processes and be able to predict properties and performance is steadily increasing.

Current State-of-the-Art Technology

Current efforts to manage nuclear weapon stockpile aging are driven largely by issues identified during stockpile surveillance. However, stockpile reduction will necessarily curtail the frequency and thoroughness of stockpile surveillance. Experiments investigating the aging of stockpile components are challenging because of the long time scales associated with aging processes, and the difficulty in determining that accelerated aging experiments provide dependable results. Current modeling efforts typically focus on specific components with suspected issues, and usually depend on large amounts of experimental input to compensate for inaccuracies in the underlying methods. In particular, the ability to predict the effects of aging through simulation is severely limited by uncertainties in the rates and mechanisms of aging processes.

The first difficulty in using current computational methods to predict aging behavior is their accuracy is insufficient to predict the rates of chemical processes under ambient conditions. Because the rate of thermally activated processes is exponential in the height of the associated reaction barrier, a small error in the barrier height becomes a large error in the rate. For example, current electronic structure methods, that can be applied to systems with hundreds to thousands of atoms, e.g., the Kohn-Sham DFT with a semi-local exchange correlation functional, predict barriers with an accuracy of roughly 5 kcal/mol, which corresponds to a factor of almost five thousand uncertainty in room temperature rates. Without further information, it is impossible to determine whether a given aging process will become significant this week or after 90 years. Reactive force fields and tight-binding methods, which allow much faster calculations, are even less accurate, with months or years of development effort required to achieve a 10 to 15 kcal/mol of across-the-board accuracy in barrier heights for a given chemical system.

Even with a method that accurately predicts reaction rates, the application of modeling to predict aging behavior would be severely limited due to the difficulty of identifying the specific processes responsible for aging. The materials involved in weapons components are usually quite complicated, involving such features as amorphous or polycrystalline structure, alloy disorder, extended and point defects, and surfaces or interfaces. The accurate representation of such materials in a quantum mechanical model requires systems containing at least 1000 atoms. Furthermore, weapons components are often exposed to a wide variety of environments; e.g., radiation, reactive chemicals, electrical discharge, mechanical stress, or temperature excursions, further expanding the set of possible aging mechanisms. In contrast, the application of electronic structure methods to determine the transition pathway for a single, well-defined dynamic process in a crystalline material (e.g., the diffusion pathway for a point defect in a compound semiconductor crystal) remains a state-of-the-art calculation requiring hundreds of calculations, each sampling hundreds of configurations. Extending such a calculation to even a simple alloy remains at the forefront of current research. In realistically complex systems, hundreds of processes are possible, and it becomes likely the key processes that actually control aging will be missed.

Science Challenges

The predictive simulation of aging requires significant improvements in the accuracy of reaction barriers relative to conventional methods, such as the DFT with a semi-local exchange correlation functional.

Considerable progress has been made in recent years toward improving the accuracy of such methods, generally at the expense of a significant increase in computational cost. For example, explicit evaluation of the nonlocal exchange has shown to improve the accuracy of calculations for a wide-range of molecular systems, but it increases the computational cost by one to two orders of magnitude. Predictive simulation of aging will require a concerted effort to further improve the accuracy of electronic structure methods with the goal of achieving "chemical accuracy," which is the ability to make realistic predictions about the rate of chemical processes at room temperature. For convenience, chemical accuracy is often taken to be less than 1 kcal/mol uncertainty in barrier heights. Quantification of these uncertainties will also be important in enabling predictive aging simulations because an atypically large error rate in one of the many competing processes can lead to inaccurate predictions from the whole process.

Further improvements in the efficiency of chemical configuration and reaction exploration will also be essential to enable predictive simulations of aging. As discussed, in realistically complex systems, there likely are hundreds of possible processes that need to be explored (each requiring hundreds of calculations and sampling hundreds of configurations) to identify the processes that actually control aging behavior. Even ignoring computational requirements and considering only the effort of the person running the simulation, shepherding tens of thousands of calculations through the computer queues by hand becomes infeasible. Furthermore, treating these simulations as independent calculations necessarily involves great duplication of effort with the same fundamental reactions being discovered over and over again. In recent years, there has been considerable progress in developing algorithms to help with the sampling of the configurations associated with a given reaction. Predictive aging simulations will require approaches that not only automatically identify the possible reactions in a system, but also string them together to identify entire pathways or even new and unexpected processes occurring over long time scales.

Need for Exascale Computing

Further improvements in electronic structure methods needed for achieving chemical accuracy will likely continue the trend of requiring a commensurate increase in computational cost. For example, evaluation of the nonlocal and long-range contributions to the correlation through a methodology such as the Random phase approximation is likely to increase the computational cost by another order of magnitude over explicit evaluation of the exchange. A calculation for a given configuration of a 1000-atom system using conventional DFT methods requires 1 teraflop-hour. A reasonable goal is to achieve chemical accuracy with the same scaling as system size, but with a factor of 1000 increase in computational requirements; i.e., 1 petaflop-hour to study 1 configuration of a 1000-atom system.

If it is necessary to study hundreds of possible processes—each requiring hundreds of calculations and each sampling hundreds of configurations in a predictive aging simulation for a realistically complex material—then 1 million configurations would have to be explored. If a chemically accurate calculation for each configuration requires 1 petaflop-hour, the full simulation would then require 40 exaflop days. Because there are hundreds of materials within a weapons system, each with its own aging behavior, it becomes clear that both significant improvements in methods and exascale computational resources are required to address this challenge.

Research Directions to Pursue

Possible research directions for improving electronic structure methods to achieve chemical accuracy are discussed in the PRD, "Speed and Accuracy of Quantum Chemistry Modeling." In addition to further development of techniques for finding transition states and accelerated dynamics methods, predictive aging simulations will require methods that can investigate very long time-scale evolution involving previously unknown pathways or processes. An example is kinetic Monte Carlo methods that automatically identify the possible transitions for each configuration of the system, while recognizing and avoiding further work for local configurations that already have been explored. A potentially fruitful approach to accelerating such calculations is to use more empirical methods, including tight-binding and reactive force field methods, as an effective preconditioner on electronic structure calculations. In this approach, fast empirical calculations would be used to identify the possible reactions for a given structure, and a much smaller number of electronic structure calculations would be used to accurately determine the barrier heights. This constant and automatic comparison of force field and electronic structure results would provide an ideal opportunity for automated improvement of force fields. Such approaches are discussed in more detail in PRD 3.

Interfacial Chemistry for Molecular Sensing and Attribution

Motivation for National Security

Predictive modeling capabilities at the chemistry-materials interface (i.e., molecules in gas or liquid solution interacting with a substrate, which could be a solid-state surface or large protein) are essential for the development of emergent sensing and attribution technologies, including device analytics. Advanced analytic methods operate at this interface, offering tremendous opportunities in nonproliferation, threat reduction, and stand-off detection, as well as chemical and nuclear fingerprinting and attribution. Device analytics include biological tethers, biomimetic motifs, and nanoarrays with emergent functionalities that, when paired with solid-state device physics, yield functionalities (olfaction on a chip, immunology on a chip, sequencing on a chip, etc.) that could provide total awareness across a broad spectrum of biological and chemical threats. Likewise, new analytic technologies such as Site-Specific Natural Isotope Fractionation-Nuclear Magnetic Resonance (SNIF-NMR) exploit complex quantum mechanical effects at the chemistry-materials interface in the detection of isotopic signatures, which may yield powerful fingerprinting and attribution capabilities related to chemical, biological, and nuclear weapons.

Current State-of-the-Art Technology

Advanced analytic methods at the interface of chemistry and the solid state are at an early stage of development and are guided almost entirely by phenomenology. There are many competing motifs. For example, current "nano-noses" involve gold nanodots on a surface, carbon nanotubes and deoxyribonucleic acid (DNA) on a surface, or olfactory proteins bound to a surface. Likewise, issues surrounding specificity, binding affinity, and signaling in biological sensors are at an early stage of characterization with detection based on a variety of native and non-native binding motifs, including aptamers, carbohydrates, native and engineered peptides, etc. Analytic technologies based on SNIF-NMR are emergent and useful for fingerprinting, but their potential for chemical attribution is at an especially early stage of development.

These early strategies are promising but have limited resolution or perfect resolution and little generality. In many cases, the design space is poorly defined. In the case of the nano-nose, the device senses a "potential" and delivers a "signal," and the burden of calibration and interpretation is placed on the software to resolve. Biological devices may have difficulty in enhancing binding affinities and signal strength. The early prototyping of SNIF-NMR for attribution also relies on software to make reliable assignments within heavily convoluted data, with little understanding of the tunneling effects that give rise to the signature.

Science Challenges

The development of engineering principles for the design and control of emergent device analytics for chemistry and biology and the rigorous interpretation of their response signature demands understanding of the underlying materials-chemistry interface at the atomistic level. In particular, three challenges for predictive modeling are apparent:

1. Acquire the ability to determine kinetic rate equations associated with quantum mechanical tunneling effects manifest in differential surface adhesion. This is related to the ability to include both electronic and nuclear quantum effects across the solid-liquid interface and the protein-liquid interface.

2. Acquire the ability to compute quantitative binding affinities for both small and large molecules, both at the solid-liquid interface and the protein-liquid interface.

3. Acquire the ability to compute changes in band structure in response to free-energy changes corresponding to surface adhesion or conformational changes in tethered macromolecules.

The challenge is not just to span multiple length and time scales, but also to extend theoretical models and computational methods across the boundaries of conventional disciplines. Many important effects at the chemistry-materials interface involve small free-energy differences that correspond to large changes in entropy and enthalpy, requiring high fidelity in both spatial and time domains. Extending time scales is difficult because, unlike materials problems with a convenient separation in the rate spectrum, these interfacial problems involve materials and liquids and biological macromolecules where many forms of accelerated molecular dynamics are inapplicable. In addition, the models and methods of classical band structure, based on local density approximations, may prove inadequate to describe the longer-range quantum effects required to model hydrogen bonding and subtle changes in band structure resulting from surface chemistry. Likewise, correlated wave-function models that capture these effects may encounter length-scale bottlenecks when applied to the solid state.

Need for Exascale Computing

Exascale computing will open the materials-chemistry frontier, allowing calculations of both enthalpy and entropy with fidelity sufficient to resolve the free-energy changes associated with differential isotope adhesion, as well as binding affinities and consequent changes to band structure. This is a premier problem at the intersection of materials science, chemistry, and biology, where the theoretical and computational difficulties are so great that exascale computing could have a significant impact.

Research Directions to Pursue

Research at the chemistry-materials interface requires breakthroughs in the treatment of novel boundary conditions, including half-slab boundaries for the solid-liquid interface, rigorous quantum-classical boundary conditions, and new unified approaches to the problem of coupled nuclear-electron dynamics. Achieving transformational results at the exascale demands more than incremental improvements in the parallelism of existing algorithms. For example, conventional quantum algorithms that scale as $O(N^3)$ with respect to system size N may experience a 1000x boost on moving from the gigascale to the exascale, while new O(N) algorithms would achieve game-changing enhancements of 10^9x. However, rigorous methods for linear-scaling electronic structure are not scalable and may require substantial software investments to implement efficiently on exascale platforms. Likewise, conventional methods for performing even classical Coulomb sums using three-dimensional Fast Fourier Transforms (FFTs) may not scale well, a problem that may be complicated by the presence of mixed boundaries and interfaces.

Predictive Chemistry of Actinides in the Environment

Motivation for National Security

One of the critical issues for DOE and NNSA is the safe and environmentally acceptable storage, processing, disposition, and management of plutonium and other radionuclides. During the Cold War, the Savannah River Site (SRS) produced plutonium and tritium in support of the nation's national defense programs. Currently, SRS has multiple missions, including environmental cleanup and restoration from years of plutonium production and production of mixed-oxide fuels for commercial nuclear power plants to dispose of plutonium as part of the Nuclear Nonproliferation Program. Radioactive liquid waste streams, produced as part of the disposition mission, are being solidified for long-term storage. Established as the Atomic Energy Commission's on-continent proving ground, the Nevada Test Site has supported more than four decades of nuclear weapons testing. One of the site's critical missions is environmental restoration, including the characterization and remediation of subsurface flow and transport of (low-level) radioactive waste sites, such as the underground test areas. Waste characterization and repackaging are conducted at the Waste Examination Facility at the site in preparation for shipment of waste for disposal at the Waste Isolation Pilot Plant in New Mexico. The NNSA complex consists of various sites where radionuclide materials are managed, stored, or disposed. For example, the Y-12 National Security Complex works to dispose of surplus highly enriched uranium in a safe, secure, and environmentally acceptable manner.

The stability of proposed radionuclide disposal sites for high-level waste, as well as radioactive waste migration through the subsurface and water table, is a difficult challenge for DOE and NNSA. A detailed understanding of the behavior, fate, and transport of radionuclides in the subsurface is needed to develop and promote safe, long-term disposal technologies for high-level waste. The reactive behavior of radionuclides in complex chemically and physically heterogeneous environments is driven by molecular-scale processes. An example is the aggregation of radionuclides, which is an important and poorly known loss mechanism that inhibits tight inventory control for transuranic elements in nuclear fuel processing (Roberto and de La Rubia 2006; Leonard et al. 1985). Ignorance about these mechanisms introduces the possibility of purposeful, chronic diversion of radioactive materials that is undetectable, yet over the course of time, can lead to accumulation of enough nuclear material for a weapon. It is also true that aggregates of radionuclides released into the environment have qualitatively faster migration rates compared to unaggregated solvated radionuclide ions (Novikov et al. 2006; Batcheller and Redden 2004).

The molecular-level mechanisms for these differences are largely unknown, and the rates cannot be predicted. As another example, plutonium may exist in four different states in environmentally relevant aqueous environments, each with a different chemical behavior. Radionuclides—such as uranium and plutonium—are known to reduce or oxidize when interacting with iron-oxide minerals in the subsurface. Mastery of the complex reaction mechanisms of radionuclide species could lead to new methods of containing radioactive releases into the environment and sensitive detection schemes for covert nuclear activity.

Plutonium Chemistry for Waste Cleanup

One of the critical issues for the National Nuclear Security Administration is the safe and environmentally acceptable storage, processing, disposition, and management of plutonium. Understanding how plutonium contaminants migrate through the environment is a major scientific challenge that currently is almost exclusively driven by experimental measurement. Computational chemistry simulations using extreme-scale computing platforms will be able to predict the complex chemistry and behavior of plutonium in dynamic real-world nuclear storage and geochemical environments (walls, long-term storage materials, minerals, and biological). Predictive simulations could enable us to control or manipulate virtually all aspects of environmental plutonium chemistry, waste processing, and separations, and thus drive scientific discovery and development of new cleanup and storage technologies while reducing the need for hazardous and very expensive exploratory experimental research.

Accurate molecular-scale models of the behavior of radionuclides in various environmental conditions provide critical information that can be used to develop more accurate field-scale radionuclide behavior and transport models for developing policies for remediation, nuclear waste processing, safety of nuclear waste repositories, and separations processes in nuclear fuel processing.

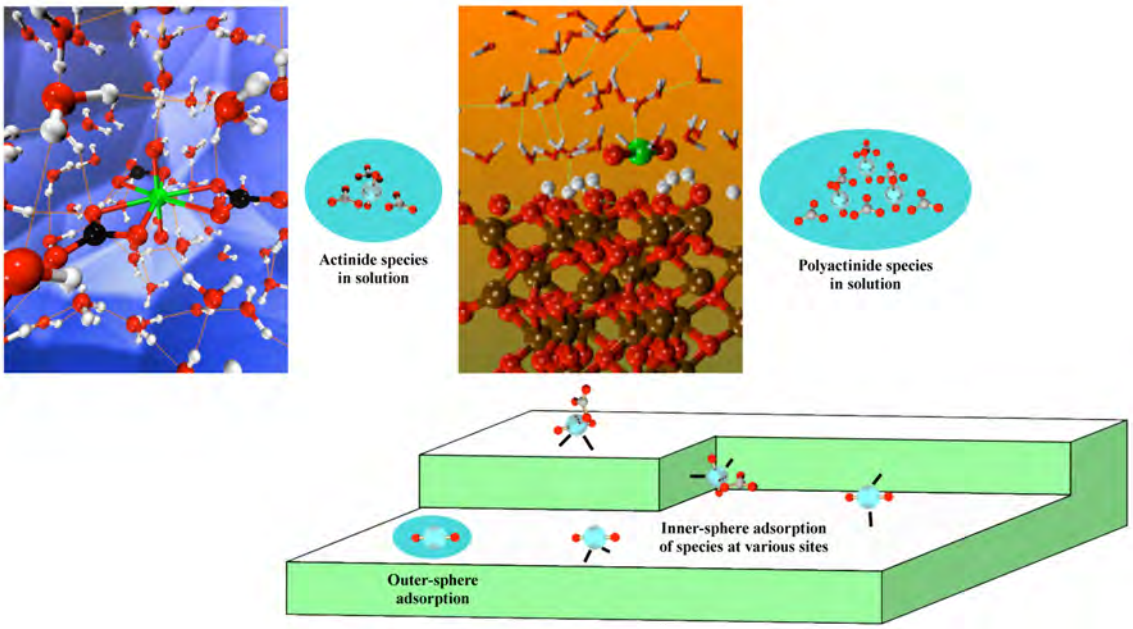

Middle image: Complex solvated plutonium species can interact with minerals in the subsurface and ions in solution through various mechanisms. Left image: *Ab initio* dynamics simulations can provide knowledge of the chemistry and stability of plutonium complexing with ions in aqueous solutions found in the subsurface. Right image: At the mineral surface, the plutonium species can be reduced, incorporated into the mineral, or new species can be formed. Images courtesy of Steve Plimpton (Sandia National Laboratories).

Current State-of-the-Art Technology

Research on the chemical behavior of radionuclides in the environment is almost exclusively driven by experimental measurement and sometimes by unexpected experimental discovery. An example of the latter is work by Soderholm et al. (2008) that led to the discovery of transuranic $Pu_{38}O_{54}40^{+}$ clusters with chemically labile and chelating outer surfaces that could be used in liquid-liquid separation of these clusters. No published computational studies have been performed on the chemistry of transuranic clusters, either of their properties in isolation or their interactions with artificial or natural (e.g., minerals and biological) environments.

With computational resources in the teraflop and petaflop range, combined with scalable software, computational chemistry studies are now starting to probe questions with reasonable accuracy related to the chemistry and oxidation behavior of radionuclide complexes at model geochemical surfaces in the presence of interacting molecules in solution. Researchers are in the early stages of modeling the structures of experimentally observed large radionuclide species, such as the clusters already mentioned. However, current computational resources limit much of this modeling to interpreting experimental data and curb the use of computational chemistry as a predictive and discovery capability for radionuclide science.

Science Challenges

The scientific challenge in this arena is for computational chemistry to become a predictive capability for modeling the complex chemistry and dynamic behavior of radionuclides in storage, processing cycles, and the environment. Predictive simulations could drive scientific discovery and development of new cleanup and storage technologies. Predictive and discovery-focused simulations will enable control or manipulation of virtually all aspects of environmental radionuclide chemistry, waste processing, and separations. This will also greatly reduce the times needed to certify these processes by more stringently bounding the uncertainties associated with these processes.

Tackling this scientific challenge requires the ability to model microscopic mechanisms in dynamic real-world process cycles, as well as geochemical and storage environments. The chemistry of radionuclides can occur under extreme conditions for parameters such as temperature, pressure, pH, or high-radiation fields. Extraction processes rely on a delicate balance of ligands, surfactants, and often high-ionic strengths. All of these factors must be included accurately to quantify these interactions, which would then aid in optimizing these processes and minimizing contaminated waste. In the environment, the actinide species could be present as isolated ions in solution. They may adhere either to solid surfaces or migrating colloids or form colloids themselves, altering their absorption and migration processes. Modeling realistic environments requires comprehensive descriptions of imperfections in materials and defects on edges on surfaces that could affect the chemistry that occurs. Again, a significant amount of solution environment (including ions) and surface chemistry must be included in the simulations to understand these processes and develop quantitative evaluations.

Beyond the microscopic scale, there is a need to account for the mechanical and macroscopic properties of materials; e.g., the formation and progression of cracks. While atomistic simulations have provided great insight into the basic mechanisms of processes, including crack formation or diffusion and phase transformations in solids, the large number of atoms in any realistic macroscopic structure currently is too large to directly simulate. Kinetic processes are critical to understanding the coupling with subsurface

flow. The size characteristics of the colloid particles, along with their associated electric double layer of counterions, can play a critical role in flow through narrow orifices where the surfaces also have charged characteristics and counterions. Consequently, large-scale simulations (thousands of atoms for electronic structure and billions of atoms for molecular dynamics) will be required.

An accurate understanding of the chemistry and behavior of radionuclides as a function of changing environmental or process conditions at the molecular level is critical to understanding basic mechanisms. This understanding can then be propagated into more accurate field-scale radionuclide behavior and transport models used to develop policies for remediation, nuclear waste processing, safety of nuclear waste repositories, and radionuclide separations processes in nuclear fuel (re)processing.

Need for Exascale Computing

Simulations of isolated radionuclide clusters or simple surface or material model systems involving ~100 atoms, of which one-third to one-half could be radionuclei, is a roughly terascale task because of the number of active electrons, the inherent relativistic effects, the presence of strong charges, and the large number of degrees of freedom. Accounting for solvent and the mobility of counterions in a realistic system is a much more difficult and approximately petascale task caused by the increase in the number of atoms and degrees of freedom, the long-range nature of ionic forces in ordering molecules and ions, and the introduction of kinetics and dynamics to explain the relevant phenomena. Modeling dynamical real-world geochemical and material environments (walls, minerals, and biological) as needed to study the phase separation of clusters in the presence of a chelation agent or the accurate description of migration rates becomes an exascale problem. In addition to more atoms and degrees of freedom, real-world simulations involve imperfect surfaces, radiation damage, and heterogeneity, all influenced by long-range ionic forces. A reliable simulation at this level of realism is an exascale challenge.

Research Directions to Pursue

Predicting the behavior of radionuclides in real-world environments will require the development of accurate computational chemistry methodologies that are fast and efficient enough to model the kinetics and dynamics of chemical processes. Dynamical reaction processes can exhibit complex behavior or rare events on time scales (from femtoseconds to seconds) that are often difficult to predict and/or identify a priori. In addition, modeling the chemistry of radionuclides in the condensed phase and at complex interfaces requires the use of methodologies that scale effectively with the system size and the development of approaches that enable the coupling of length scales ranging from atoms to macroscopic systems.

PRIORITY RESEARCH DIRECTIONS

Advancing the Speed and Accuracy of First-Principles Chemistry

Scientific and Computational Challenge

There are three main methods for the *ab initio* prediction of the electronic of molecules and materials: 1) DFT, 2) wave-function methods, and 3) quantum Monte Carlo (QMC). There are also several emerging methods that will be summarized in this panel report. Each approach has its merits, and this situation is expected to remain for the foreseeable future, thus requiring continued development of all methods.

For small molecular systems, wave-function methods have repeatedly demonstrated their ability to attain chemical accuracy for systems near equilibrium, including bond energies (~0.1 kcal/mol), bond lengths (~0.001 angstroms), and vibrational frequencies (~1 cm^{-1}). For light elements, such accuracy already requires the inclusion of relativistic effects though perturbation theory. For heavy elements, it is necessary to incorporate fully both scalar and spin-orbit relativistic potentials. The central idea is the construction of a hierarchy of *ab initio* many-body methods that incorporate increasing levels of excitation from a reference model and employ increasingly large one-particle bases. From both theory and comparison with experiment, it is known these approaches converge to the exact result. The same methods also apply for excited states and away from equilibrium, but the calculations are correspondingly more complex and expensive. The increased complexity is due to the simultaneous need to describe multiple electronic states and to treat multiple unpaired electrons that can couple in a myriad of ways. The latter issue is referred to as the "open-shell problem." Scaling of computational cost with both molecule size and accuracy limits the applicability of these methods to larger molecules (discussed in more detail as follows).

QMC methods can obtain similar accuracies for small molecules and their cost increases less rapidly than wave-function methods with the size of the molecule. However, the unresolved problem of enforcing Fermi-Dirac statistics without introducing approximations makes it difficult to systematically improve results, and many properties—especially those of excited states—are not yet accessible via QMC.

Due to its lower computational complexity, DFT is the method of choice for routine computation on molecules of any size. The exact functional is unknown, and the most successful functionals are semi-empirical. In combination with a variety of empirical corrections, DFT can attain near chemical accuracy for many ground-state energies. However, away from equilibrium, these corrections do not apply, and systematic defects in the approximate functionals lead to a complete inability to describe specific types of excited states, electrical and optical properties, and certain types of weak interactions. Ignorance of the exact functional implies there is no systematic way to improve results.

Summary of Research Directions

The scientific challenges of the preceding section have been divided into three parts: 1) solving the electronic structure problem, 2) solving the chemical problem, and 3) making contact with experiment. As noted, the electronic structure problem can be solved to any desired precision (in principle), but current approaches are impractical for large systems. This is due to the highly nonlinear scaling of computational cost with the size of the system and desired accuracy. For instance, the successful Coupled-Cluster with Single and Double and Perturbative Triple excitations method increases in cost as the seventh power of the system size and as the fourth power of the accuracy (i.e., doubling the size of the system increases the cost by a factor of 128 and requesting an extra decimal place of accuracy increases the cost by 10,000-fold). Density functional methods typically scale as the third power of system size. Physically, an optimal complexity nearly linear in system size and logarithmic in accuracy, even quadratic scaling, would be a significant advancement. Presently, there are various reduced-scaling methods that unfortunately all suffer from various defects, particularly uncontrolled, size-dependent errors that can be introduced in an effort to force linear scaling.

Assuming the electronic structure problem can be solved for any configuration of the atoms, solving the chemical problem of predicting chemical energies and reaction rates requires the incorporation of the effects of atomic motion to compute vibrational contributions and partition functions. Except at very high

temperatures, it is necessary to incorporate quantum effects, and the partition function is exponentially sensitive to the treatment of low-frequency modes that typically are strongly anharmonic. For floppy molecules comprising circa 30 atoms, the limiting factor in predicting reliable reaction rates resides not in solving the electronic structure problem, but in the vibrational problem. New theories and computational tools are required to advance beyond the harmonic or classical approximations.

Finally, making contact with experiment requires computing observables of models that incorporate all pertinent complexities of the real system. Direct computation of observables (e.g., Raman or X-ray spectra) rather than inferred quantities (band gaps) yields much more information, eliminates assumptions, and provides greater confidence in successful comparison. This need increases the emphasis that must be placed on computing not only ground state, but excited state and response properties. In addition, experimental samples systems are often poorly characterized and inhomogeneous. There are various theoretical approaches for handling this uncertainty; e.g., those used to model metal alloys, but none are understood in the context of accurate wave-function methods.

Even armed with an optimal computational method, exascale computing is required for many NNSA chemistry applications. The expected asymptotic near-linear scaling of an optimal method will only be obtained for systems that exceed a problem-dependent, minimal size. Below this size, the computational cost typically increases nonlinearly because of the physics involved. The prefactor of the linear term can be understood to be the cost of computing inside this minimal volume. Incorporating atomic motion, either through explicit dynamics or computation of anharmonic force fields, is hugely expensive and presently is only possible at the terascale/petascale for modestly sized systems at low levels of theory. Full treatment of relativistic effects using two-component spinors (four-component methods probably are not necessary) increases the computational cost at least another 16-fold.

Without a focused and concerted effort to develop a new generation of theoretical methods and associated computational tools, the necessary computational chemistry capabilities will not be deployed on the time scale demanded by NNSA's pressing strategic and environmental goals, nor in time to realize the full potential of exascale computers. The current pace of progress is too slow to deliver the necessary theoretical breakthroughs over the next decade, and the cost of a new generation of codes far exceeds current levels of support. Moreover, for computational chemistry to mature beyond being primarily a tool for basic engineering and scientific discovery, new elements to manage uncertainty in the simulations must be completely incorporated and embrace the incomplete knowledge and uncertainty inherent to experiment and the real world. In consequence, these robust and predictive tools will enable exascale chemical simulations to inform decision-makers and serve as true design tools.

Suggested research topics include the following:

- **Reduced-scaling many-body methods.** Develop consistent theoretical/algorithmic frameworks that allow the optimal scaling to emerge from the physics rather than as a result of ad hoc modifications. New basis sets, including explicitly correlated methods, seem to be essential for a breakthrough in this area, and novel approaches—such as reduced density-matrix and Green’s function methods—have the potential for breakthrough. NNSA’s interest in heavy metals and reactive species requires progress also be made in the open-shell problem.

- **Improved density and current functional.** Develop new functionals that better treat multicomponent, reacting systems, long-range forces, excited states, and explicit time evolution.

- **Solving the atomic motion problem.** Develop effective methods to incorporate the vibration and dynamics of quantum nuclei on anharmonic surfaces into chemical energy differences and partition functions.

- **A new generation of codes.** To manage the cost and complexity of software associated with modern theory and the anticipated architectures of exascale computers, a new generation of codes is required. These codes must enable and encourage rapid deployment of novel theories and algorithms, be able to address reactive multiphase and multispecies systems, and be scalable to hitherto unimagined levels of parallelism. A multidisciplinary team model patterned after DOE's Scientific Discovery through Advanced Computing (SciDAC) Program team seems the most appropriate for initial development transitioning to long-term management as a facility.

Expected Computational and Scientific Outcomes

"The right answer for the right reasons"[6] is a central mantra of modern computational chemistry that, over the last two decades, has yielded practical tools for all elements in the periodic table, which are capable of quantitatively accurate predictions of many chemical properties. Robustness and accuracy are the intrinsic value of first principles, or *ab initio* methods, over the much faster (semi-)empirical methods. Work on the tasks outlined in this PRD will provide a substantial improvement in both metrics.

"Robustness" means the sought accuracy is guaranteed (regardless of chemical composition, displacement away from equilibrium, external fields, etc.) because all of the physics are resolved. Robustness is especially critical for many NNSA applications due to an emphasis on reactive, multispecies, nonequilibrium systems exemplified by interactions and chemical transformations in high-energy and extreme environments. "Accuracy" means that—in principle—the exact result is obtainable, and there is a hierarchy of methods of increasing computational complexity that systematically yield improved accuracy. In contrast, (semi-)empirical methods require results from both experiment and first-principles simulation before they can, if ever, be extended into new regimes. Such methods also offer no systematic approach to improved fidelity.

Potential Impact on National Security

All of the theory and methods discussed in this PRD apply broadly to quantum chemistry applications; specifically to actinide chemistry. The next PRD, "Predictive Actinide Chemistry from First Principles" focuses on the special complexity issues raised by actinide compounds and their unique role in the NNSA mission.

Predictive Actinide Chemistry from First Principles

Scientific and Computational Challenges

Prediction of the properties of molecules and materials containing actinide atoms is one area of crucial importance to DOE's missions. Actinide atoms occupy the last row of the periodic table and include uranium and plutonium. Their properties are determined by their electronic structures, which are complicated because these are atoms whose valence electrons are in the so-called "f-shell." It is

[6] In chemistry, this phrase and its implied emphasis on connecting simulation to rigorous theory is attributed to Professor Ernest Davidson.

important to have accurate ways of predicting the properties of compounds containing actinide atoms because, in many cases, experiments are difficult to conduct because of safety concerns or the materials themselves are difficult or expensive to make.

Summary of Research Direction

The first step in delivering accurate predictions for the properties of actinide molecules and materials is the calculation of their electronic structure. This requires solving the Schröedinger equation, including all the electrons and nuclei in the system, not just those in the actinide atoms. For moderately complex compounds, this number could require many thousands of electrons and even hundreds of thousands to millions for more complex systems. Current technology for the prediction of the electronic structure of actinide systems solves an equivalent formulation to the Schröedinger equation called DFT. State-of-the-art programs are capable of routinely solving for the electronic structure of systems with only a few hundred electrons or for very long calculations up to about 1000.

Although DFT has a 40-year history, further advances are still required. DFT is based on a description of the electron-electron interactions dominated by so-called exchange-correlation functionals. Exchange-correlation functionals used today are largely split between hybrid functionals in chemistry and gradient functionals in solid-state physics. Unfortunately, these functionals are not yet able to treat mixed molecular and solid systems with the chemical accuracy (1 kcal/mol) needed for accurate predictions. To make matters worse, in the area of molecular solids where van der Waals interactions are important, there is no consensus on best practice. Moreover, the state of these models for the physics and chemistry of f-electron elements and compounds also is unsatisfying because traditional functionals can—and do—fail for even basic properties.

Hence, fundamental work in the development of accurate electronic structure techniques is needed, specifically for f-electron materials. This will require formulation of more accurate exchange-correlation functionals. Because of the large system sizes (numbers of electrons), new numerical and parallel algorithmic advances will be required to run these systems efficiently on exascale platforms.

Once the electronic structure of stationary actinide compounds can be predicted, the next step will be to enable modeling of their dynamics, such as when they undergo chemical reactions. This will require a combination of electronic structure methods with molecular dynamics. The latter requires the calculation of the forces on all the nuclei, which are determined by the interactions between the electrons, as well as by the electrostatic repulsions between the nuclei. The cost of these calculations, performed for many steps at the appropriate time scale, will require an exascale level of computing resources.

Expected Computational and Scientific Outcomes

To date, combining electronic structure calculations with molecular dynamics has been possible only in relatively simple cases. Hybrid functionals are the most powerful for describing molecular processes today, but they require at least an order of magnitude larger computational speed than gradient-based functionals. Being able to do long, large-scale molecular dynamics simulations with hybrid functionals would increase confidence in first-principles simulations for an abundance of chemistry applications, including dissociation and reactions under pressure, ions in liquids, catalysis in dense environments, etc. Having the option to routinely do first-principles molecular dynamics with hybrid functionals will require extreme computing, but it can significantly improve overall understanding in the short term.

Potential Impact on National Security

With dramatically increased computational resources, the physics and chemistry of actinides is an area likely to reach a tipping point. There are several technical reasons for the difficulties in modeling actinides: the large number of electrons, the mixed core/valence character of f-electrons, relativistic effects, and shortcomings of existing density functionals. To enable the state of the art for actinide physics and chemistry to reach a level comparable to that achieved for first-row elements and compounds, (close to) all-electron molecular dynamics, in combination with new functionals, are needed. Actinide chemistry/physics is an area of general interest to NNSA and timely given the growing interest in nuclear energy and in nuclear weapon surety without testing.

Reactive Molecular Dynamics with Quantified Accuracy

Scientific and Computational Challenges

The ability to perform atomistic-scale simulations that accurately capture the physics and chemistry in materials and molecules is of vital importance for a number of areas closely related to the NNSA missions. Examples include aging of nuclear stockpile materials (BSC 3), highly explosives sensitivity and initiation (BSC 2), and the interfacial chemistry of sensors (BSC 4). All of these areas require dynamic simulations on multicomponent systems at multiple length scales. While exascale computing can increase the time and size-range restrictions of first-principle quantum mechanical simulations to the nanometer/nanosecond level, realistic simulations of the events related to aging, initiation, and interfacial events will require simulations at the micrometer/microsecond scales. To accomplish this for the diversity of these systems with successful sampling of the relevant physics and chemistry will require a substantial improvement in length and time scales covered by current reactive molecular-dynamics simulations. Accurate simulations at these scales will enable reactive molecular dynamics (RMD) to become predictive (rather than explanatory), interact directly with experiment, and guide manufacturing decisions.

Simulation methods that are a couple of orders of magnitude faster than quantum mechanical methods are needed to facilitate super-microscale simulations. This speed improvement can be obtained by replacing the physical rigor of quantum mechanical methods with a level of empiricism. These empirical methods must display the following two attributes:

- Reliability: Provide quantifiable levels of chemical accuracy for ground states, material properties, and reaction barriers.

- Transferability: Enable simulations on systems with a large chemical and structural diversity, ideally covering the entire periodic table.

Currently, empirical methods that show promise to fulfill these criteria include tight-binding approaches and more empirical IP. However, substantial additional development is required for these methods before they can occupy the microscale niche on exascale computer systems. Likely, this development will lead to more complex empirical schemes. This trend has been most apparent in the IP field, where elegant and compact early models (Lennard-Jones potential; Tersoff potential) have given way to more accurate, but more complex, potentials (e.g., MEAM [Lee and Baskes 2000], ReaxFF [van Duin et al. 2001], and AIREBO [Stuart et al. 2000]). This increase in complexity and parameter space is unavoidable and has

been sustained by the availability of larger computational resources and larger, quantum mechanically derived training sets. The research directions outlined as follows aim to provide the framework for the development of empirical methods suitable for the exascale era.

Summary of Research Direction

Development of Self-Checking, Self-Aware Empirical Methods

By projecting the trends in empirical methods reliability and transferability demands to the exascale level, it is clear the complexity of empirical methods will become too substantial to manage using chemical and physical intuition alone. As more phenomena are added to the model, the opportunities for unforeseen consequences proliferate. Potentials that are accurate under a wide range of different chemical environments cannot be built by hand—only through exhaustive testing of many different atomic configurations and consideration of many different candidate potentials. This will require the development of advanced sampling tools, where predictions from the empirical methods will be compared to quantum mechanical results and, if required, added to the empirical methods training set. Furthermore, new empirical methods optimization tools, including evolutionary optimization methods (genetic algorithms, genetic programming), will be required to allow for use of large training sets and extreme computing to search for high-accuracy "needles" in the IP haystack. An example of successful generation of a three-body potential via genetic algorithm sampling is shown in Figure 17.

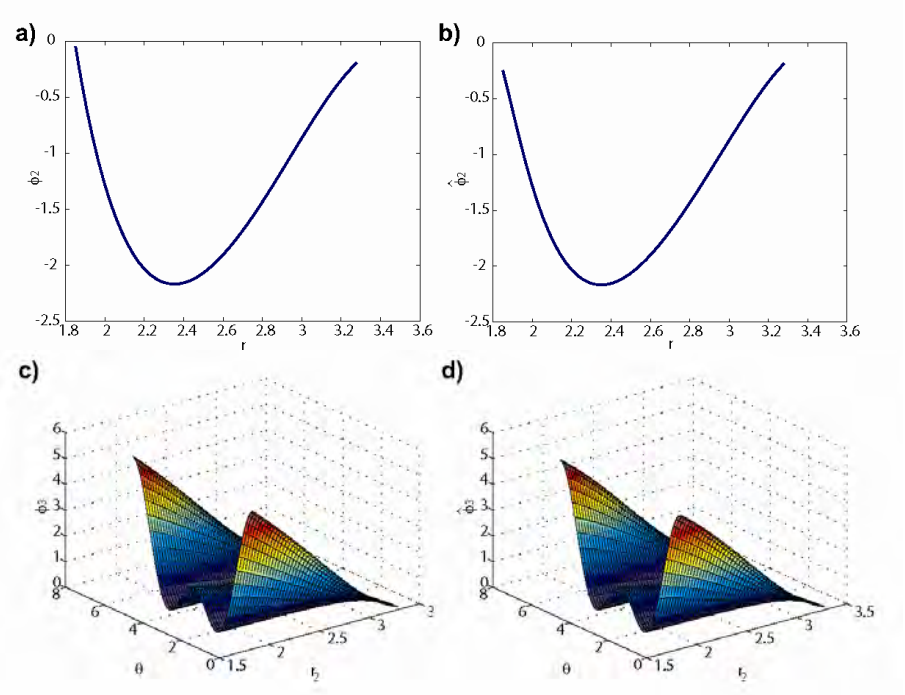

Figure 17. Potential energy surface for the Stillinger-Weber three-body potential for silicon, as a function of bond angle (theta) and distance (r_2). The left plot is the analytic S-W expression. The right plot is from a genetic algorithm that effectively generated the same energy dependence on theta and r_2 without knowledge of the formula. Image courtesy of Michael Brown (Sandia National Laboratories).

A successful coupling scheme between empirical methods and quantum mechanical methods should allow the empirical methods to become self-aware, enabling local error estimates that give the empirical methods authority to interrupt a simulation to initiate a series of quantum mechanical calculations in regions of maximum uncertainty, and to incorporate the results of these quantum mechanical data into its representation of the potential surface. In the small molecule world, the above strategy has been used to automatically grow spectroscopically accurate (e.g., ~cm-1) or kinetically accurate (e.g., ~1 kcal/mol) potential energy surfaces (Dawes et al. 2009; Camden et al. 2009; Netzloff 2006). In most recent applications, the expensive quantum mechanical calculations needed by the empirical method fitting methods have been distributed over tens of processors. The scaling of such methods to much larger systems of national security interest and their incorporation into simulations primarily driven by IP representations is a vital area of research.

Development of Schemes for Empirical Method Functional Form Optimization

Going beyond parameter optimization schemes, the optimization tools should also be able to suggest and employ modifications in the empirical method functional representations. Two examples of loosely coupled parallel global optimization strategies are replica exchange Monte Carlo and hierarchical fair competition, both of which have been used to optimize IP using genetic programming (Slepoy et al. 2007; Brown et al. 2010). In both cases, an ensemble of populations of individual candidates is constructed. These populations are evolved in isolation from each other, save for infrequent exchanges between neighboring populations. By using different search strategies in different populations, it is possible to search a large space while still preserving the fittest individuals. By combining the ability for empirical methods to optimize both their functional forms and parameter space, scientists can obtain self-aware, self-checking empirical method models that will enable predictive large-scale atomistic simulations of systems critical to NNSA missions.

Corrections of Nonclassical Effects in Empirical Method-Based Molecular-Dynamics Simulations

Even when empirical methods are developed that fulfill the reliability and transferability criteria, the fundamental issues with classical molecular dynamics approaches must be addressed. A significant number of materials of interest have light atoms and stiff bonds that lead to normal modes with high vibrational frequencies whose dynamics are not well approximated by the classical equations of motion used in molecular dynamics simulations. Under these circumstances, regardless of the accuracy of the IP used to calculate atomic forces, the resulting dynamics will be incorrect and could lead to inaccurate predictions. This problem arises when the quantum of energy for the oscillator is larger than the available thermal energy. Classically, a harmonic oscillator can have any amount of energy, but quantum mechanics restricts the possible energies to integers of the quantum of energy. Significant differences between the two descriptions arise when only a few excitations are present.

While equilibrium properties calculated from molecular dynamics simulations can be corrected to incorporate quantum effects, their role in nonequilibrium dynamical processes (e.g., chemical reactions) is not yet fully understood and remains one of the fundamental challenges in condensed matter physics and chemistry. Examples of important applications where quantum dynamical effects are expected to significantly influence molecular dynamics predictions are as follows: decomposition, reaction and detonation of molecular explosives, radiation damage of semiconductors, aging, and electron transfer processes. Exascale computing opens the possibility of applying path-integral methods (Reichman et al.

2000), bohmian dynamics (Gindensperger et al. 2000), or alternative approaches to large-scale reactive molecular dynamics with accurate IP. A successful effort will significantly increase the predictive power of atomistic simulations of chemical reactions in condensed phases.

Expected Computational and Scientific Outcomes

The development of a verifiable and transferable computational method for atomistic simulations at the super-microscale will allow atomistic-scale simulation tools to become predictive, enabling targeted design at the nanoscale. This will radically modify the role of computational science. It will enable direct communication with experiment and will allow atomistic simulations to guide—rather than follow—experiment.

Potential Impact on National Security

Exascale computing will enable quantifiable, accurate atomistic simulations at a size range that allows consideration of the full complexity of the materials. Such simulations will provide accurate predictions of critical material interface phenomena related to stockpile aging, high-energy material initiation, sensor chemistry, and environmental restoration. These simulations will afford a more complete description of complex materials than experiment, providing an accurate evaluation of the state and reliability of devices critical to national security.

Seamless Multiscale Chemical Reactivity

Scientific and Computational Challenges

Many national security problems involve chemical reactivity that spans multiple length and/or time scales. For example, a metal part could corrode due to oxidizer migration over a time of several years. The state of the corroded part could vary on scales from centimeters to nanometers. This is a heterogeneous, multiphase problem. The national security challenge is to predict the behavior of the corroded part under extremes of pressure in excess of 10^6 atmospheres and temperatures in excess of 10^4 Kelvin. Clearly, this class of problems challenges traditional chemical simulations in several ways. Traditional approaches take either a continuum reactive flow approach (Veynante and Vervisch 2002) or an atomistic molecular dynamics approach (van Duin et al. 2001). Usually, these approaches are specialized to a narrow range of thermodynamic conditions. Future success in predicting multiscale chemical reactivity will require close cooperation between atomistic and continuum simulation.

Exascale computing platforms expand the possibility of simultaneously studying realistic problems at the same problem size from both the continuum and atomistic perspectives. Current molecular simulations are limited to roughly 10^9 atoms. Similarly, current continuum simulations are limited to roughly 10^9 zones. An exascale computer could treat 10^{12} atoms on a routine basis. This is sufficient to study a material representative volume element with a length of 10^{-6} m, which achieves overlap with mesh sizes in highly resolved continuum simulations. In the future, hybrid approaches, which combine elements of atomistic and continuum simulation, are expected to become a viable path to understanding complex multiphase reactive systems.

Summary of Research Direction

The issue of long time-scale simulations requires a combination of algorithmic breakthroughs and exascale computational resources. Most current approaches to parallel atomistic simulation have increased the system simulation size without increasing the time scale of the simulation. In most cases, it is difficult to exceed time scales in excess of 100 nanoseconds in molecular dynamics simulation with current methods. Most quantum simulations of reactivity are currently limited to about 10 picoseconds, while sophisticated reactive force fields typically are limited to 100 picoseconds. Many interesting problems in chemistry, however, involve time scales that greatly exceed these limits. While traditional approaches such as transition-state methods are viable for slow chemical reactivity in gases, the coupling between the reactive site and the environment makes these methods difficult to apply to complex sequences of reactions in condensed phases. Replica-based approaches to long time-scale simulation are promising (Voter 1998), but still have limitations that must be overcome before they can be applied to multistep chemistry. Exascale computers will enable massive replica dynamics to sample rare events. Communication between replicas could potentially be used in the development of more highly efficient long time-scale methods.

Expected Computational and Scientific Outcomes

The ability to predict chemical processes across a wide range of time and length scales would be a revolution in chemistry modeling that could apply to diverse areas, such as biology, materials science, and chemistry, and span basic and applied science. Predicting the aging of materials is of particular importance for national security, although it also has obvious application to manufacturing. Development in this area could lead to explosive materials with dramatically improved safety characteristics.

Potential Impact on National Security

Although long-term work is likely necessary to treat reactive problems on arbitrary length and time scales, the overlap of atomistic and continuum simulations at the exascale could affect national security issues within 10 years.

CONCLUSIONS

Table 4 lists the milestones for the work described in this section. Provided that the computational resources become available for research at the anticipated scales, the forefront research activities that will be conducted are provided as milestones.

Table 4. Milestones for Chemistry Panel

Scale	Milestone
>1 Petaflop-year	• Reactive molecular dynamics (RMD) of a $(0.1\ \mu m)^3$ sample for 1 μsec • High explosive (HE) modeling of reaction zone in single crystals • Quantum chemistry (QC) of 2500 atom system for 100 psec • QC modeling of phase transitions in heterogeneous bulk system
>20 Petaflop-years	• RMD of a $(0.3\ \mu m)^3$ sample for 1 μsec • HE modeling of reaction zone affecting voids (single crystal) • QC of 5000 atom system for 100 psec • QC modeling of phase transitions at extended interfaces
>100 Petaflop-years	• RMD of a 1-μm^3 sample for 1 μsec • HE modeling of reaction zone in polycrystalline sample • QC of 10,000 atom system for 100 psec • QC modeling of 1-part-in-1000 inertial confinement fusion (ICF) mixtures
>1 Exaflop-year	• RMD of a 1-μm^3 sample for 10 μsec • HE modeling of initiation & deflagration transitions (polycrystalline) • QC of 25,000 atom system for 100 psec • QC modeling of phase transitions versus composition for lo-ratio alloys/mixtures • QC modeling of equation of state (equation of state) for HE product gases/solids
RMD baseline: 10^4 flops per atom per time step, 0.1 femtosecond time step, O(N) scaling. QC baseline: Vienna Ab-initio Simulation Package (VASP) time (500 light atoms, 30 MD steps, 8 Tflops) = one hour, 0.5 femtosecond time step, $O(N^3)$ scaling.	

THE SCIENCE OF NONPROLIFERATION

Co-Leads: Mark D. Rintoul, Sandia National Laboratories
Alex Pothen, Purdue University

Panel Members: Mihai Anitescu, Argonne National Laboratory; Edward Van Eeckhout, Los Alamos National Laboratory; Christopher Oehmen, Pacific Northwest National Laboratory; and Jaideep Srivastava, University of Minnesota

CURRENT STATUS

Nuclear nonproliferation refers to the activity of preventing the spread of nuclear weapons both directly and indirectly through control of the nuclear material and technologies required to make a usable weapon. Nonproliferation is one of the most complex activities that the U.S. Government engages in due to the significant political and technological difficulties associated with the spread of nuclear technologies. Establishing nonproliferation policy is difficult because almost all aspects of the nuclear weapons development process have a complementary peaceful process associated with nuclear energy production or general scientific research.

The current administration has made nuclear nonproliferation a key component of its nuclear strategy. In April 2009 in Prague, President Obama delivered a speech to the international community to emphasize the United States' renewed commitment to "to seek the peace and security of a world without nuclear weapons."[7] This speech outlined the political steps necessary to put the world on the path towards a future without nuclear weapons. The speech focused on reducing the current nuclear stockpile, strengthening the nuclear nonproliferation treaty, and building a framework for civil nuclear power cooperation. However, the speech was also forthright in identifying that there will be a significant challenge in the implementation of such a path, stating that "… we go forward with no illusions. Some countries will break the rules."

Given the political difficulties in enforcing a policy of nuclear nonproliferation, it is necessary to have a clear technological advantage over those countries whose leaders are tempted to break the rules. Such an advantage is not only important in detecting offenders but serves an even more valuable role in deterrence. If the United States can make covert proliferation activities economically infeasible for other nations to engage in the practice, then diplomats' jobs become that much simpler.

Current Nonproliferation Strategy

Although the issue of nonproliferation reaches across many cabinet agencies, there are two that do much of the work. The U.S. Department of State, through its Bureau of International Security and Nonproliferation, primarily works on policy issues. The U.S. Department of Energy (DOE), through its Nuclear Nonproliferation organization within the National Nuclear Security Administration (NNSA), focuses on many of the operational technological issues. However, there is significant overlap and interaction between these and other agencies to comprise the total national nonproliferation effort.

[7] April 5, 2009. "Remarks By President Barack Obama." Hradcany Square, Prague, Czech Republic. Entire speech is located at http://www.whitehouse.gov/the_press_office/Remarks-By-President-Barack-Obama-In-Prague-As-Delivered/.

The Department of State's nonproliferation work covers all types of nonproliferation (including nuclear, chemical, biological, and other), but a significant fraction of that work applies to nuclear nonproliferation. The agency has key responsibilities for developing the policies related to the international nuclear community including activities around the Nuclear Non-Proliferation Treaty, nuclear export controls, and the physical protection of nuclear facilities. In addition, the agency assists other nations with their own nonproliferation activities with advice and strategic funding opportunities. Finally, the Department of State helps develop diplomatic responses to detected nuclear activities.

The nonproliferation activities of the DOE can roughly be categorized into three areas: detect, secure, and dispose. Detection is completed in multiple environments. On the ground, it consists of deploying radiation systems at strategic border crossings, airports, and seaports. In the air, it consists of specialized detectors looking for evidence of nonproliferation or detonation activity. DOE's securing mission emphasizes keeping both weapons and nuclear material away from state and non-state actors by helping other countries secure the existing weapons-related materials and helping peaceful nuclear activities continue while minimizing the proliferation risk. Finally, disposal remains a critical mission of DOE. Plutonium continues to be produced by many nuclear reactors in the world as part of the energy-generation process, and this remains a dangerous source of potential weapons-grade material along with the nuclear material that is taken from weapons in the disarmament process.

Crosscutting Nonproliferation Science Topics

The science of nonproliferation is not an independent scientific discipline; rather, it is a term used to describe the collections of sciences and technologies needed to enable the nonproliferation mission. Although the focus of this section will primarily be on nuclear nonproliferation, many of the general topics are applicable to other types of nonproliferation issues including potential biological and chemical weapons. The different sciences that are relevant to nuclear nonproliferation are extremely diverse, primarily due to the very different challenges of detecting and preventing proliferation at different points in the weapons-production cycle. These different scientific disciplines are generally well studied for their importance outside of the realm of nonproliferation, but this section will highlight a few of these different disciplines and emphasize the particular aspects that are relevant to the nuclear nonproliferation science.

Nuclear Physics

The fundamental underlying science surrounding effective nuclear nonproliferation is nuclear physics; there are several reasons for this. First is the need to understand the precise scientific and engineering concerns surrounding the nuclear enrichment process. This understanding is needed to predict what steps an adversary must take to produce weapons, whether from external procurement or internal development. The second reason that nuclear physics is important is related to nuclear detection. This is somewhat more specific to the nonproliferation mission. The focus on detecting different nuclear materials is extremely important to nonproliferation science and is less important to other applications of nuclear physics such as nuclear power plant design. The science around detection generally lies at the interface of nuclear and solid-state physics, because the challenging problems in detection are related to the interaction of nuclear decay by-products (neutrons, gammas, antineutrinos, etc.) and the material used to make the detector. While it is not necessary for the nonproliferation science community to fund much of the work in fundamental nuclear science, it is generally important for that community to help support the development of better detectors.

Remote Sensing

Although remote sensing is not generally considered a science itself, it is a fairly mature discipline that has a long history of research and impact. Although remote sensing is formally defined simply as sensing at a distance, this section primarily uses this term to refer to imaging across the electromagnetic spectrum (especially visible and nearby wavelengths) via aircraft and satellites. The critical science areas behind remote sensing are generally in the area of physics on the collection side of the imaging process and in the area of computer science on the analysis side of the process. Remote sensing is critical to many civilian and military applications and has a robust technology base that nonproliferation can tap into. However, many of the algorithmic challenges associated with automated analysis of images have somewhat unique nonproliferation issues. Increasingly, automated analysis and data extraction from images will be important to the field.

Information Science

Nonproliferation science is not so much a "formula-driven" science as it is a "data-driven" science. Because of this, information science techniques are becoming increasingly important in the field, especially as the relevant data sets (which include images, text data, and sensor data) become larger. There are two primary parts to this problem. The first is related to how the data are extracted and stored, and the second is related to how the data are accessed and explored. The first is somewhat nontrivial, especially with respect to data storage. There are many disparate types of data that need to be stored in a way that they can effectively be used as a whole. Many Internet search engines are struggling with this problem as image, video, and audio data proliferate on a media that was once dominated by text. Of course, the text-mining problem has been studied for many years within the information science community, and it too remains a problem that has only been partially solved for specific applications. The question of accessing and exploring the different types of data needed to make decisions is also complicated by the different types of data.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

The panel studied the issues that face the science of nonproliferation and found that most of the challenges broadly fell into three major categories: 1) improve the effectiveness of physical devices in proliferation detection; 2) aggregate large and diverse data corpora; and 3) empower experts to explore and analyze aggregated data.

Improve the Effectiveness of Physical Devices in Proliferation Detection

Much of the limitation in identifying nonproliferation lies in physical detection abilities. For ground-based radiation detection, this includes gamma and neutron detection. In the case of remote sensing, this includes better observational capabilities for satellites. Both of these can be significantly helped by advances in materials science and manufacturing to build better physical devices. However, improving the physical devices is just one part of advancing the effectiveness of proliferation detection. In general, significant progress remains to be made in analyzing the data from physical devices. In most cases, the data contain uncertainties associated with but not limited to device imperfections, fundamental physical limitations, and statistical noise. Figure 18 demonstrates how a better mathematical understanding of the uncertainties associated with the detection model would result in more precise devices and enable a more quantitative understanding of the risks associated with decision-making based on detector evidence.

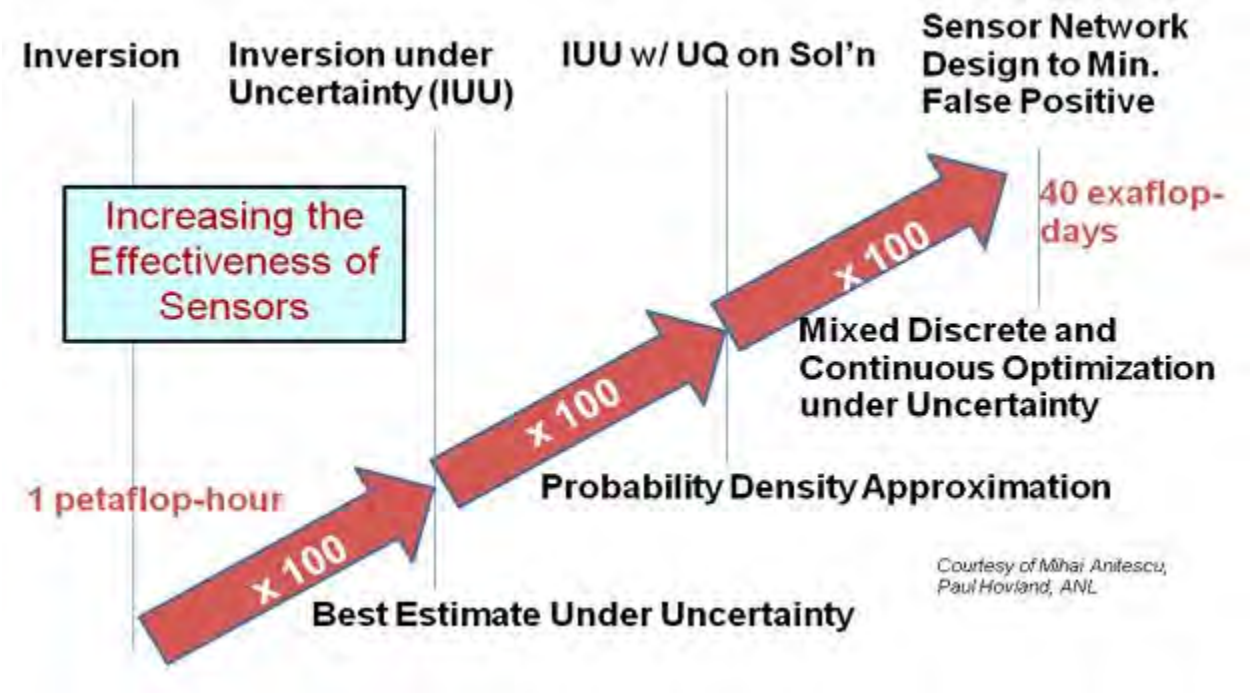

Figure 18. An example of the increasing need for computational power to enable revolutionary improvements in nonproliferation sensing. Image courtesy of Mihai Anitescu (Argonne National Laboratory) and Paul Hovland (Argonne National Laboratory).

Aggregate Large and Diverse Data Corpora

Researchers are moving toward an era with an increasingly large amount of data being generated in nonproliferation activities. This includes spectral data associated with sensors, image data, text data that have been generated through many sources, and even observations of human activity. All of these data must eventually be aggregated such that they can be optimally used in modeling and analysis. Bringing together disparate data for effective use remains a difficult computer science problem for nonproliferation and for the broader Internet community. There are currently simple, ad hoc means of doing this for specialized situations, but understanding how to do this in a general sense for abstract data types remains a distant goal.

Empower Experts to Explore and Analyze Aggregated Data

Ultimately, the effectiveness of the collected data is only as useful as the information that an analyst can extract from it. After the different types of data have been collected, techniques for an analyst to explore the data will be needed. This can involve specific types of queries where an analyst must have a means of phrasing a complex question of the data and have the computer return one or more potential answers and the associated confidence with the answer. Perhaps more powerful is enabling an analyst to explore the data in a more nonspecific way to look for potential patterns or anomalies.

PRIORITY RESEARCH DIRECTIONS

Optimization, Inversion, and Design

Optimization is the rational process by which one selects the best among all alternatives that conform to a set of predefined specifications. Optimization problems are defined by their variables (the parameters that represent the alternatives) and their key data functions; the objective, which defines the criteria ranking the alternatives; and the constraints, which define the specifications. The flexibility of these simple modeling principles makes optimization one of the most powerful computational paradigms.

Examples of such applications include the following:

- Inverse problems. One example is a source inversion problem. In this case, the variables are the source distribution. The objective function could potentially be the least-squares discrepancy between the measured signature and the one obtained by the propagation of the proposed source distribution through the forward physical process. Finally, the constraints would be the non-negativity of the flux or other prior information.

- Sensor design problems. This is where the variables include the physical design and technological choices in the sensor, the objective function is a mix of its performance criteria, and the constraints are the manufacturing and deployment specifications. These problems include asset management in the case of a threat response where the objective function can be time to respond, and the constraints are the resource or transportation limitations.

Nevertheless, while optimization as problem formulation is a common interface between a variety of problems in nonproliferation, the performance of optimization algorithms is intimately tied to the exploitation of the structure of the problems solved. It is therefore worthwhile to investigate both optimization as a general paradigm for nonproliferation science, as well as specific distinguished instances that occupy a large portion of the nonproliferation activities such as inverse problems and sensor design.

Scientific and Computational Challenges

Nonproliferation science problems will involve integer and continuous variables, nonsmooth data functions, and imperfect prior information. Therefore, current optimization techniques must be adapted and extended to deal with these new challenges.

To solve problems that have integer and continuous variables, recent techniques of solving mixed integer nonlinear programs must be customized and improved. Such approaches start with a branch-and-bound strategy where scenario branching is done on the various values of the integer variables. This is then combined with a cutting-plane approach, where inequality constraints are added that separate the region where the optimal solution can be proved not to lie from the region that is still searched for the optimal solution. In many recent approaches, combinations of these two strategies form the basis for defeating the perceived exponential complexity that integer variables add to the problem. To deal with lack of smoothness in the data functions, recent research has shown substantial improvement can be made by incremental smoothing combined with additional inequality constraints. Nonsmoothness can appear, for instance, as a result of the modeling of the action options of the adversaries in game theoretical

approaches. Investigation is needed to determine how problems of the size and complexity generated by realistic applications can be resolved. Finally, a "silver lining" of the applications in nonproliferation is the lack of complete information about the problem to be solved. Bi-level or robust optimization approaches can be used to determine worst-case scenarios. In some circumstances, however, it is appropriate to pursue a probabilistic approach. Such is the case where the threat is moderate while the safeguarding actions can incur considerable cost. The appropriate paradigm for deciding alternatives in this circumstance is the one of optimization under uncertainty. Probably the best-known variant of this approach is stochastic programming where either the objective function or the constraints contains statistics—such as mean, variance, or risk measures—of the various outcomes. Exploring scenarios in a large uncertainty space is one of the most difficult problems in computational science and carries with it the characterization of the "curse of dimensionality." In optimization under probabilistic uncertainty, one must study techniques to reduce the variability of the outcome that is due to the fact that complexity prevents the inclusion of all scenarios in the optimization approach. The hierarchical nature of uncertainty—combined with a superior understanding of the interplay between optimization and uncertainty by means of, for example, polynomial chaos expansions—can be exploited to reduce the variability to levels appropriate for decision-making.

Summary of Research Direction

Numerical optimization has made significant progress for the case where its variables are continuous and its objective and constrained functions are smooth, with the exception of problems with bound constraints (range constraints on an individual variable), and highly structured problems. Problems with billions of variables are commonly solved on several thousand processors. Important advances have also been observed in problems where the variables are discrete (such as 0 to 1 variables that quantify, for instance, the availability of a discrete asset, such as a field detector). Here, the combinatorial nature of the problem is an obstacle to a rapid increase in the size of the variable set resolved. Nevertheless, algorithmic advances have allowed the resolution of problems of high theoretical complexity far faster than could have been expected from hardware progress alone.

Therefore, a lesson learned from recent optimization advances is that modeling decisions should weigh in the expected performance of the algorithms. This points to the benefits of integrated teams of nonproliferation and optimization solver experts and to the important role of algorithmic advances in high-reward computational capability development.

Expected Computational and Scientific Outcomes

Given the recent progress in large-scale optimization and related high-performance computing (HPC) technologies, one candidate problem for early demonstration on exascale architectures is the one of source inversion, an inverse problem. Exascale computing is particularly important when the solution of the inverse problem must fully account for the effects of incomplete information and is used in developing the design of the security response itself. Indeed, recent advances in inverse problems in computational geosciences, which share many features with source inversion problems, have shown that problems on a regional scale can be solved on the order of 1 petaflop-hour. The addition of a representation of the uncertainty in the function data, a probabilistic solution of the problem, and the use of the inverse problem in a decision setting with mixed integer and continuous variables are each expected to increase the computational complexity by factors of roughly 100. As a result, making a consequential decision based on the best available inverse-problems technology with full awareness of the

risk would require on the order of 40 exascale-days of computational effort. This assumes optimal embedding of existing HPC technology, such as multigrid linear algebra, automatic differentiation (for the computation of the tangent linear models needed to accelerate the sampling), and mixed integer nonlinear programming technology on this never-before-tested scale. Researchers' current theoretical understanding, as well as the experience acquired at the petascale, suggests this is achievable in the next decade with exascale computing.

Potential Impact on National Security

Researchers' abilities to perform inversions on sensor data would remarkably reduce the uncertainty in the results of their sensors. This is significant in areas where the resolution of their sensors cannot be improved because of physical or geographical constraints. Using existing technology to produce raw data with less uncertainty would significantly decrease the risk in very high-consequence decisions made on sensitive nonproliferation issues.

Proliferation Process Modeling

Scientific and Computational Challenges

Model building and validation for proliferation will require the use of large and diverse data sets, given all the possible permutations of processes throughout the entire cycle. These data sets will be large and will require the exploitation of uncertainties throughout any realistic analysis. This extends the problem to an extreme framework, particularly if Monte Carlo analyses are invoked.

Summary of Research Direction

The processing, enrichment, and casting of uranium and plutonium into weapons are often done separately at different geographic locations and at different times. Thus, it can be very problematic to tie each disparate facility (performing individual or multiple processes) into an overall picture of proliferation. In addition, because each process can often be performed in multiple ways, the observations or signatures from each could vary markedly. To simulate the many probabilistic perturbations of these processes is truly daunting.

Thus, models of the proliferation process that combine probabilistic linkages among entities, such as finance, shipping, physical security, procurement of materials, infrastructure, feedstock, processing, effluents, and final products, need to be created. Processes such as mining and milling, ore processing, enrichment, reactor fuel fabrication, high explosives and electronics fabrication, high explosive testing, tritium extraction, reprocessing, parts fabrication, pit creation, nuclear testing and assembly must be included.

Information about these operations is obtained in several forms ranging from small, discrete packets of data to large quantities of seemingly random or disconnected data sets. Furthermore, conceptual semantic graphs and ontological models are required to extract information from these data sets because discerning useful information often requires linkages between disjointed, transient, and typically heterogeneous inputs. These ontological models will clarify the connections among multidisciplinary sources.

Expected Computational and Scientific Outcomes

By modeling and simulating the processes that might lead to proliferation, one will be able to start unraveling the complexity of what multiple signatures or observations might indicate. The probabilities of certain indicators might really stand out as "silver bullets." New computational advances are expected to be realized in parallel, handle uncertainties in the data and model, and combine disparate data types.

The quantification of uncertainty through this simulation would be most helpful in bounding the proliferation problem. Uncertainties that will need to be dealt with are the model itself, the strength of various measurements within the model (i.e., will the measurement affect the conclusions directly or in some indirect manner?), and evidence uncertainty—particularly for disparate data. Because data can be from various sources (e.g., human intelligence sources with high, medium, and low confidence; chemical sources with a mean and standard deviation; or visual sources such as photographs), it is challenging to quantify and use the uncertainties of each.

To perform this simulation with the assistance of extreme computing would allow many more scenarios than might be possible under current capabilities. For example, Monte Carlo simulation, which would be an appropriate methodology for assessing the uncertainties involved, would require extensive parallelization and further algorithmic construction than is currently available.

Potential Impact on National Security

Clearly, proliferation process modeling is very important to understanding the status of nuclear capabilities within certain countries. These simulations will help the U.S. Government establish clear guidelines for establishing the level of confidence of their observations. That is, how confident is the Government that certain facilities might or might not be engaged in proliferation, given what is known?

The entire problem set has not been formalized because it is very large and complicated. Some modeling efforts for individual facilities exist today—but if extreme computing can be used, probabilistic assessments throughout the entire proliferation process modeling cycle can be assessed for individual or multiple facilities.

Information Extraction and Aggregation

Scientific and Computational Challenges

To better understand what is meant by "information," researchers will use a paradigm that is becoming increasingly accepted by the research community. In this paradigm, there is a hierarchy of understanding associated with any sort of facts: Data→Information→Knowledge→Wisdom. Roughly, these terms are defined in the following:

- **data**, which are simply the set of symbols, both qualitative and quantitative
- **information**, which is data that have been processed such that specific relationships between the data are called out

- **knowledge**, which is information that has been processed in a way that it can be used and is obtained by calling out patterns in the information and data

- **wisdom**, which is the application of integrated knowledge to draw extrapolative conclusions regarding future nondeterministic events.

This is obviously a very coarse semantic description of a subtle set of concepts, and there are different ways to define the same terms and different ways to decompose the hierarchy of concepts. However, for the sake of simply describing the contents of this section, these definitions will suffice. The key challenges of information extraction and aggregation are finding computational means of turning data into information in a way that it can be used by humans and machines to obtain knowledge (and ultimately wisdom).

Summary of Research Direction

The focus of this priority research direction is effectively on the conversion of data to information. The processes of extraction and aggregation are called out separately because they represent a well-defined way to describe the steps necessary for the conversion of data to information. In addition, this priority research direction will also describe the specific computer hardware and architecture issues that the extraction and aggregation problems identify. Turning information into knowledge is closely aligned with the fourth priority research direction, which is focused on information exploration and will be discussed in the next section.

Expected Computational and Scientific Outcomes

A variety of data types are used to build information about nonproliferation. In some cases, such as spectra collected from sensors, the relationship between the data and what information it represents is relatively straightforward. In other cases, the data must be significantly processed to extract information. One of the most well known cases is image data. A large number of overhead images have been collected and are constantly being collected by a number of different means. Image data are an excellent example of pure "data," because they are simply a collection of bits (in most cases) that are information free until, in most cases, examined by a human analyst. Before this happens, the bits are converted to a set of colored and spatially structured pixels that allows an analyst to pick specific patterns and spatial relationships that represent information. This process of identifying the patterns and spatial relationships in images has long been the target of an enormous amount of research in computer science, and this remains one of the areas that HPC could affect enormously—not only for nonproliferation, but for many other areas of science and society.

There are two primary reasons why using computation to extract information from images could be a game-changing technology. The first is that annotating images is currently a rate-limiting step in many types of analyses because there is most often a human in the loop. Progress in solving the computer-image analysis problem would allow the current process to go much faster and could enable a completely new paradigm in imaging where a much larger number of images could be obtained and analyzed. The second driver for computational image analysis is somewhat more revolutionary and is very tightly coupled to the notion of "data to information." There is a strong need to capture the information in images in a way that can be used for computers to turn that information into knowledge. This means not only trying to identify features in an image, but to capture the list of features and the relationships

between them in a way that can be meaningfully used in a computational (or at least a computationally assisted) way. In general, this could lead to a reduction by many orders of magnitude in storage costs if users did not need to store the images themselves but the information contained in them.

Another important source of data is text, especially in electronic form. Most publicly available text data are available through the Internet, and even private data usually have an electronic representation that can be easily used. Again, text is data and not information until the information relevant to the need at hand can be extracted. Textual analysis and summarization can also be a "bottleneck" if completed by a human analyst. Natural-language processing is a rapidly maturing field in computer science that deals with extracting information from text. This can be completed with a specific analytical end in mind or in a more discovery-focused way through which anomalous events can be found. Human-behavior observational data are a related type of data. In this case, these data are often recorded as a text description and must be summarized in a way that can be used by computers.

For definition purposes, "aggregation" means the process of storing the information extracted from the data and finding a way to represent the information and the relationships between the different pieces of information. In computer science, the computational representation of a set of related concepts is called an ontology. An ontology is critical to storing and using much of the information that is collected in a non-numerical form. A nonproliferation ontology is currently under study within the NNSA's NA-221 Office of Proliferation Detection, and several large-scale computational issues have emerged.

One significant issue with data represented by a nonproliferation ontology is related to the uncertainty associated with the information. Whether the data have a rigorous scientific uncertainty associated with errors in measurement, or a less well-characterized certainty that is associated with human speculation, the storage and manipulation of the uncertainty of the data is a major computational challenge. Another major issue related to information aggregation has to do with the large number of different types of data and their inter-relationships. In many cases, there are conceptual relationships (works at, is a friend of, etc.), geospatial relationships (is south of, is 10 miles from, etc.), and temporal relationships (visited 5 years ago, always happens before, etc.). The combinatorial explosion of the different types of relationships makes storage and manipulation of such entities difficult.

Expected Computational and Scientific Outcomes

The amount of information that could be gathered using a more computational approach to information extraction would be orders of magnitude larger than currently possible and could enable a significant revolution in analysis capabilities. Computer-assisted filtering and organization of information would enable human analysts to focus on only the highly relevant pieces of information, and to have access to a more complete description of a nonproliferation scenario.

Information Exploration

Scientific and Computational Challenges

The goal of nonproliferation sciences is to accurately identify entities engaged in acquiring the means to develop nuclear capabilities or testing, distributing, or acquiring nuclear devices or components. Nonproliferation science is one of many scientific fields that is increasingly dominated by the availability and volume of data from disparate sources including sensors, modeling and simulation, open-source

multimedia, and conventional numerical calculations. Extracting actionable knowledge from this rich data landscape is a significant challenge because events of interest often hide in a vast sea of benign activity. Many of the methods for finding these events are exploratory in nature because the underlying behaviors they seek are not governed by canonical equations and so cannot be discovered by solving equations alone.

Information exploration is an iterative, user-driven process in which users interact with data to discover, develop, and refine hypotheses from the data themselves through sorting, filtering, projection, classification, clustering, and a host of other mathematical techniques. Unfortunately, most users currently have to make the unpleasant choice between throwing away data because they are too voluminous—a process that will prevent one from finding the unexpected—or accepting "heroic" computing times (months to years) as the rate-limiting step in the analysis process. Neither option is acceptable in nonproliferation sciences where the results of conclusions, which drive the severity of actions, require extreme fidelity; also, acceptability is affected by the time sensitivity of arriving at an actionable conclusion.

Extreme-scale computing is therefore poised to become an essential component for emerging nonproliferation sciences because it allows for data-intensive computing algorithms and techniques that rapidly drive sorting, filtering, classification, and other numerical tasks to inform meaningful visualizations and data representations in real time or near-real time without throwing away data. However, realizing this vision of HPC-enabled analysis requires many technical, algorithmic, and theoretical advances and may drive requirements for unconventional hardware, operating systems, or other components due to the scale of the problem (Figure 19). The goal is to advance proliferation sciences to enable rapid and accurate data analysis from many different sources of enormous volumes through extreme-scale algorithms-, hardware-, and systems-research programs. The remainder of this section describes how these advances can be realized by focused research programs that ultimately will drive transformational capability developments in nonproliferation sciences.

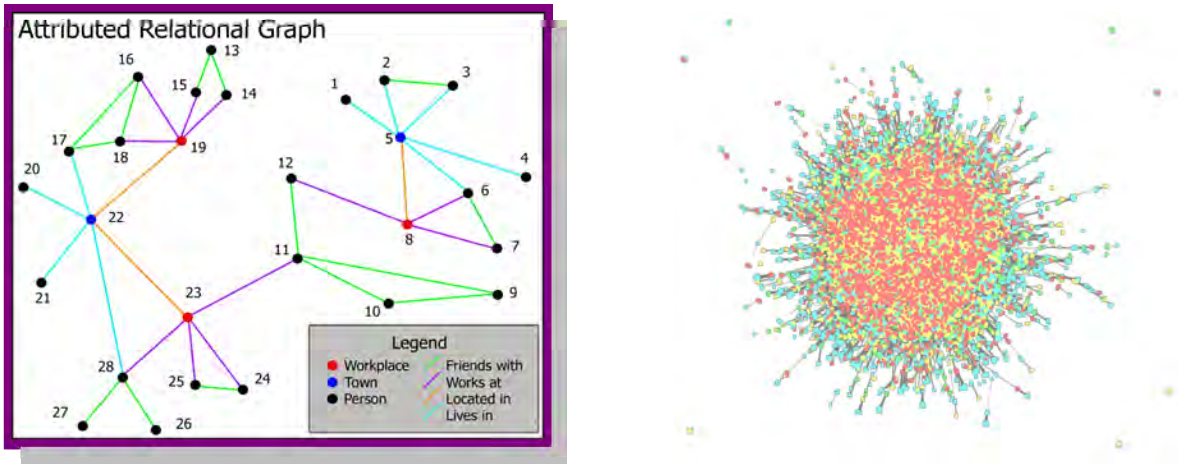

Figure 19. An example of the difficulties in exploring information: Can one find the first pattern in the second? Image courtesy of Tammy Kolda (Sandia National Laboratories).

Vision and Roadmap of Extreme-Scale Nonproliferation Sciences

Researchers envision that extreme-scale computing will play a significant role in nonproliferation sciences in the foreseeable future. One reason for this is that data volumes and rates will probably continue to grow exponentially, exacerbating the already overwhelmed analysis systems in place. While advances in hardware are part of the answer, hardware improvements alone are not sufficient to address the data problem. In addition, no single architecture will be suitable for all the needs of nonproliferation sciences at this scale. Proliferation sciences will likely need to take advantage of advances in extreme-scale computing across the full spectrum of systems including heterogeneous, distributed architectures like compute clouds or grids; more tightly coupled collocated, distributed memory systems like large-scale clusters; and highly coupled large-scale shared memory systems—all of which would incorporate some combination of multithreading and multicore processing.

To make the most of this array of analytical horsepower for the highly data-intensive application spaces of nonproliferation sciences, algorithms, middleware, visualization, and modeling and simulation capabilities will need to converge, likely leveraging a variety of different architectures in the same workflow. The game-changing advances needed to enable researchers to extract knowledge from these enormous data resources will need to enable users to ask high-level questions that cannot be answered using a single data set or analysis technique. Examples of how this integrated approach might look to an end user are described in the following scenarios.

Scenario 1: Real-time, iterative hypothesis discovery. Many researchers are tasked with extracting actionable knowledge from data sets. As nonproliferation data continue to grow exponentially with increased sensor output rates, computing capacity, and data availability, researchers most often have to discard what is believed to be "unimportant" to allow computational tractability. In addition, they often have to start with a hypothesis and then use the data sets to find supporting evidence. One problem with this approach is that with large enough data sets, one can find evidence to support any conclusion. A second problem is that there may be crucial, unexpected knowledge in a data set that is missed because it is not related to the hypothesis under investigation. Extreme-scale computing can change the paradigm of knowledge discovery from one of hypothesis testing to hypothesis discovery. Combining extreme-scale computing infrastructure using user-directed middleware (to create analytical pipelines, join processes, automatically move data between analysis modules, etc.), one can drive analysis at a large enough throughput to operate on all—or a much larger fraction—of the available data. Using visual metaphors that are browsable, analysts can get a "sense" of what is happening to devise a hypothesis from the data itself. Leveraging extreme-scale compute capacity would dramatically reduce the analysis time so that rapid, accurate, actionable (or at least, testable) knowledge may be extracted from large-scale datasets.

Scenario 2: Integrating data from multiple "sources." Signal analysis, modeling, and simulation have been mainstays of nonproliferation science from its inception. As models grow in sophistication, many have begun to incorporate more human elements (e.g., intent, readiness, etc.). Getting these values correct can directly affect the fidelity of a model—but these values can be nearly impossible to infer because one must extract them from a sea of natural-language text, images, and nonspectral signals (in addition to spectral or hyperspectral signals). While extreme-scale computing may play a direct role in signal analysis, modeling, and simulation, it may also play a critical role in improving these calculations by informing them with knowledge extracted from the enormous, open-source multimedia data also available to researchers. However, informing these models and calculations using multimedia sources

does not fit into the conventional computing model. Hybrid analysis pipelines that span heterogeneous and homogeneous clusters and tightly coupled shared memory platforms will be necessary.

Expected Computational and Scientific Outcomes

Realizing these scenarios and others enabled by extreme-scale computing will require development of the following:

- algorithms for operating on large-scale data sets such as clustering, signal, or signature detection; machine learning and classification; graph/network construction; and analysis

- specialized hardware implementations of some of these algorithms using, for example, general purpose graphics processing units, field-programmable gate arrays, highly multithreaded or multicore (or combinations thereof) platforms, heterogeneous platforms like compute clouds or grids, high-performance clusters, or large-scale shared memory platforms

- middleware that allows researchers to join analysis tasks on different platforms, to move and track data and analysis tasks, and to drive large-scale visualization

- visual metaphors for describing relationships between data elements and building a higher-level understanding of these relationships.

Potential impact on National Security

DOE has a long-standing body of research in nonproliferation sciences that would provide a sound basis for national security at the extreme scale. DOE is uniquely positioned to integrate extreme-scale computing into nonproliferation sciences because of its leadership role across the HPC spectrum. DOE is home to mature applications for data management, integration, and workflow development as well as visualization. There will likely be a strong mathematical research element to algorithms that operate on data at the extreme scale. Applications will have to deal with missing data, corrupt data, uncertainty quantification and aggregation, data fusion, statistical issues that arise with extremely large data sets, and other challenges. These research areas already have many mature applications within the DOE research portfolio. It is expected that similar expertise to that which has advanced basic and applied sciences can be brought to bear on proliferation sciences.

Successfully enabling nonproliferation sciences at the extreme scale will enable researchers to rapidly find the unexpected and transform the security landscape from hypothesis driven to data driven (where the best hypothesis is discovered rather than guessed at). This can be done by data-driven analysis informed by the complete body of data available to researchers and decision-makers by integrating extreme-scale computing, algorithmic advances, sophisticated visual metaphors, and browsing capabilities in integrated workflows.

Statistics and Machine Learning for Detecting Rare and Anomalous Behavior

Large, complex, distributed systems generate huge volumes of multimedia, multiresolution data in applications such as avionics, manufacturing systems, transportation systems, computer networks, etc. Rare events in such data are events that occur very infrequently (with frequencies from roughly 5% to less

than 0.1%). For example, in the cybersecurity domain, the number of network connections that are attacks is an infinitesimal fraction of the set of all connections—and in the credit-card domain, the number of fraudulent transactions is similarly a very tiny fraction. Similarly, in vehicle health management (for aircrafts, trucks, etc.), the numbers of events that have (adverse) health impacts are few. Although rare events are by definition infrequent, in each of these examples, their importance is quite high compared to other events—making their detection extremely important.

Anomaly detection refers to the problem of finding patterns in data that do not conform to expected behavior. These nonconforming patterns are often referred to as anomalies, outliers, discordant observations, exceptions, aberrations, surprises, peculiarities, or contaminants in different application domains. Of these, anomalies and outliers are two terms used most commonly in the context of anomaly detection, sometimes interchangeably. Anomaly detection finds extensive use in a wide variety of applications such as fraud detection for credit cards, insurance or health care, intrusion detection for cybersecurity, fault detection in safety-critical systems, and military surveillance for national defense activities. The importance of anomaly detection is because anomalies in data translate to significant (and often critical) actionable information in a wide variety of application domains. For example, an anomalous traffic pattern in a computer network could mean that a hacked computer is sending sensitive data to an unauthorized destination. An anomalous magnetic resonance image may indicate the presence of malignant tumors. Anomalies in credit card transaction data could indicate credit card or identity theft, or anomalous readings from a spacecraft sensor could signify a fault in some component of the spacecraft.

Background Literature

There has been some work in the past on the detection of rare events and anomalous behavior. Traditionally, the difficulty has been in having sufficient examples from which rare events and anomalous behavior can be characterized. However, with the large amounts of data collected today, there has been a dramatic increase in the level of interest. Researchers from the machine-learning/model-building community are developing novel techniques and applying them to a number of problems. Weiss (2004), Kumar et al. (2005), and Chadola et al. (2009) provide background information on those topics.

Research Problems

Over time, a variety of rare-event detection and anomaly-detection techniques have been developed in several research communities (Figure 20). Many of these techniques were specifically developed for certain application domains, while others are more generic. There are several promising directions for further research in anomaly detection. These directions include the following:

- Detection techniques for anomalies, as well as contextual and collective rare events, are beginning to find increasing applicability in several domains, and there is a large scope for development of new techniques in this area.

- The presence of data across different distributed locations has motivated the need for distributed-rare-event and anomaly-detection techniques. While such techniques process information available at multiple sites, they often have to simultaneously protect the information present at each site, thereby requiring privacy-preserving anomaly-detection techniques.

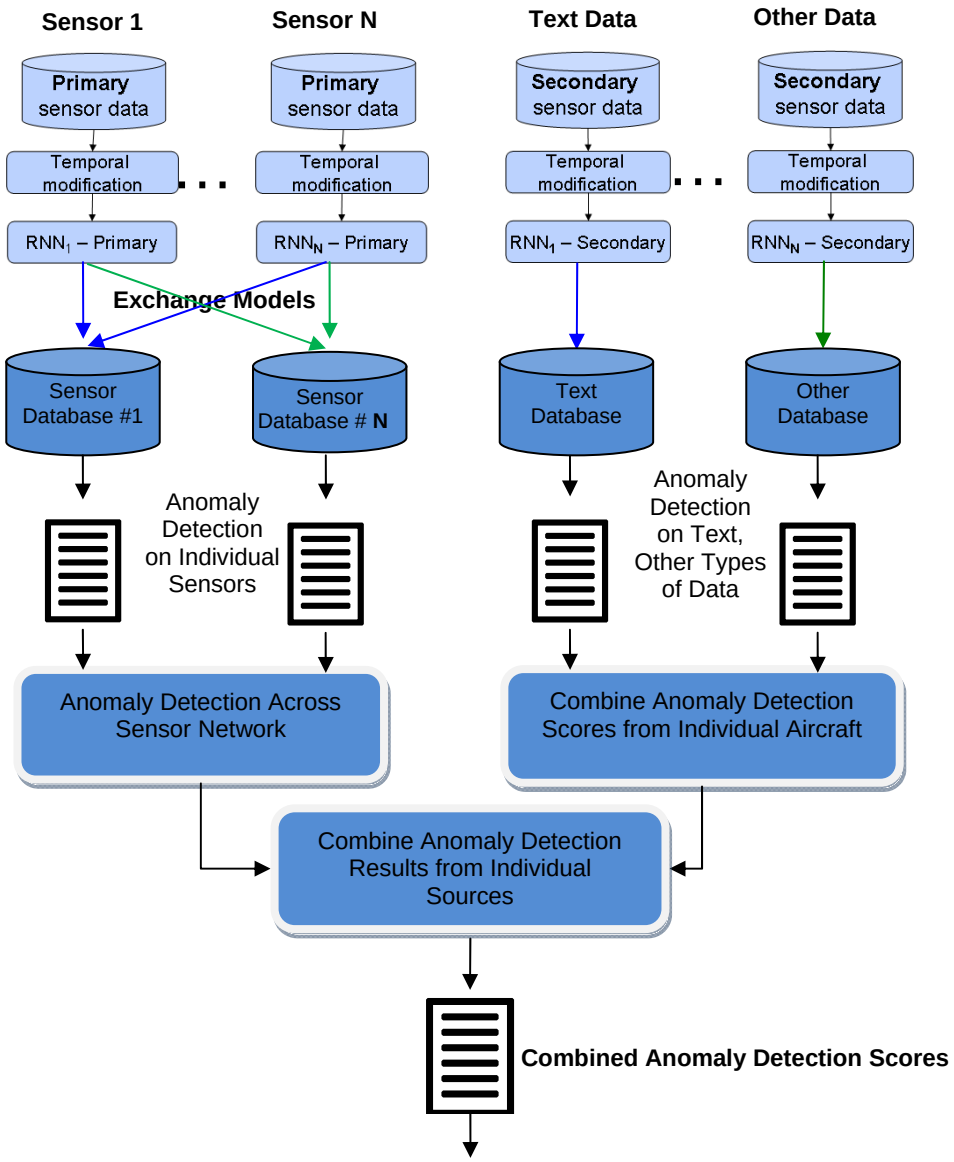

Figure 20. A generic framework for rare-class and anomaly detection. Image courtesy of Jaideep Srivastava (University of Minnesota).

- With the emergence of sensor networks, processing data as they arrive has become a necessity. Many techniques require the entire test data before detecting rare events and anomalies, but that may no longer be viable in many settings.

- Recently, techniques have been proposed that can operate in an online fashion. Such techniques assign a score for a rare event and/or anomaly to a test instance as it arrives, and they incrementally update the model.

- Another upcoming area where rare-event and anomaly detection is finding more applicability is in complex systems. An example of such a system is an aircraft system with multiple components. Rare-event and anomaly detection in such systems involves modeling the interaction between various

components.

- Real-world data tend to show highly multimodal behavior, and different techniques might work best for each mode of behavior; this needs to be considered.
- Finally, ensemble methods have shown great promise in statistics and machine learning, and their applicability in this domain needs to be explored.

CONCLUSIONS

Nonproliferation science will be a critical driver for exascale-level computing in the future. The extrapolation to exascale computing is easily seen in the extension of existing mathematical techniques to levels that require more data, more precision, and a better understanding of uncertainty. However, it is even more important to note that nonproliferation science will drive completely new types of computing. The research and development required for exascale computing in nonproliferation science will require a focus on faster computers and on new types of algorithms, data storage and retrieval, data interaction, and computing architectures. This combination of activities will have an impact in the area of nonproliferation and will allow DOE to affect the future of general-purpose computing in a fundamental way.

UNCERTAINTY QUANTIFICATION AND ERROR ANALYSIS

Co-Leads: Dave Higdon, Los Alamos National Laboratory
Richard Klein, University of California, Berkeley, and Lawrence Livermore National Laboratory

Panel Members: Mark Anderson, Los Alamos National Laboratory; Mark Berliner, Ohio State University; Curt Covey, Lawrence Livermore National Laboratory; Omar Ghattas, University of Texas; Carlo Graziani, University of Chicago; Salman Habib, Los Alamos National Laboratory; Mark Seager, Lawrence Livermore National Laboratory; Joseph Sefcik, Lawrence Livermore National Laboratory; Philip Stark, University of California, Berkeley; and James Stewart, Sandia National Laboratories

> The uncertainty is as important a part of the result as the estimate itself…An estimate without a standard error is practically meaningless. H. Jeffreys (1967)

Uncertainty quantification (UQ) studies all sources of error and uncertainty, including the following: systematic and stochastic measurement error; ignorance; limitations of theoretical models; limitations of numerical representations of those models; limitations of the accuracy and reliability of computations, approximations, and algorithms; and human error. A more precise definition is UQ is the end-to-end study of the reliability of scientific inferences.

Ideally, UQ results in the following:

- a quantitative assessment of that reliability
- an inventory of possible sources of error and uncertainty in the inferences and predictions
- an inventory of the sources of error and uncertainty accounted for in the assessment
- an inventory of assumptions on which the assessment is based.

UQ for estimation, prediction, and assessment has long been held as fundamental to scientific investigations. Traditionally, UQ has been conducted via statistical analyses in applications ranging from drug efficacy trials to inferring the speed of light. Such analyses typically rely on a mix of theory, basic mathematical models, and sufficient observational or experimental data.

Advances in computing over the past few decades—both in availability and power—have led to an explosion in computational models available for simulating a wide variety of complex physical (and social) systems. These complex models—which may involve millions of lines of code, and require extreme-computing resources—have led to numerous scientific discoveries and advances. This is because these models allow simulation of physical processes in environments and conditions that are difficult or even impossible to access experimentally. However, scientists' abilities to quantify uncertainties in these model-based predictions lag well behind their abilities to produce these computational models. This is largely because such simulation-based scientific investigations present a set of challenges that is not present in traditional investigations. These challenges include the following:

- The amount of physical data (observational or experimental) is typically quite limited.
- The computational demands of the model limit the number of simulations that can be conducted.
- The computational models are not perfect representations of physical reality—they have inadequacies, approximations, missing physics, etc.
- The computational models typically have unknown parameters and boundary conditions that need to be adjusted for the application at hand.
- Researchers often wish to extrapolate such models to conditions where they have little or no physical observations to validate model output.

The sidebar on the following page provides a very simple example of how experimental observations and a computational model are combined to infer times of flight for an object being dropped from a tower. In addition to leading to more accurate predictions, advances in UQ methodology will lead to more reliable uncertainties for simulation-based predictions. This is particularly important in high-consequence decisions for which both understating and overstating uncertainties leads to excessive costs or liabilities. More importantly, new UQ methods for simulation-based investigations will lead to improved understanding of the different sources of uncertainty affecting predictions. This will allow decision-makers to be more effective with their limited resources. For example, one may ask, "How can I best use my available resources to reduce uncertainties? Should I improve computing resources? Should I carry out new experiments? Which ones? Should I improve experimental diagnostics? Should I improve existing computational models?"

CURRENT STATUS

Presently there is a substantial amount of research activity devoted to inference and UQ aided by computational models. Some of the main areas of current research are outlined below. These current research areas relate to, and help motivate, the priority research directions (PRDs) presented at the end of this section.

Inverse Problems and Calibration of Computational Models

A computational model requires physical observations to adjust key model parameters, initial conditions, and/or boundary conditions to better model the physical system. In a typical inverse problem, these quantities are determined by minimizing the discrepancy between physical observations and computational model output. Some statistical approaches to inverse and calibration problems require that this discrepancy between observations be formalized into a likelihood function, which is produced from a probability model for the data given the model parameters. More formal statistical inference about the unknown parameters and/or initial and boundary conditions can then be made to describe their uncertainty. Because many of these inverse problems are ill-posed—especially when estimating a large field of initial conditions—many approaches regularize (enforce "smoothness" or other properties on) the unknowns being estimated. Bayesian methods for statistical inversion and calibration have become popular because these approaches codify regularization in prior distributions and give a probabilistic description of the resulting uncertainty.

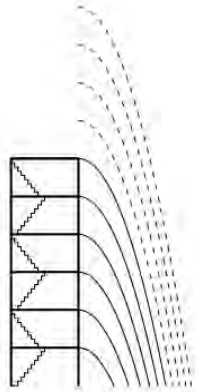

A simple UQ example:
Using experimental data and a computational model to predict drop times from new heights.

Experimental data:
The time it takes an object to drop from each of 6 floors of a tower is recoreded. There is an uncertainty in the measured drop times of about $\pm.2$ seconds. Predictions for times are desired for drops from floors 7 through 10 – which do not (yet) exist.

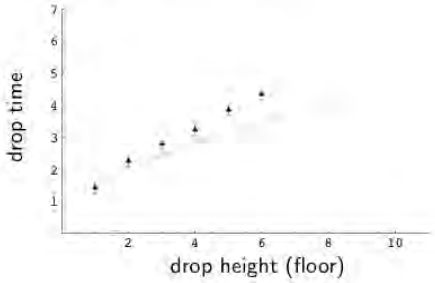

Simulated drop times:
A (trivial) computational model is developed to predict the drop times as a function of drop height. The simulated drop times (red line) are systematically too low when compared to the experimental data (plotting symbols). The error bars around the observed drop times show the observation uncertainty.

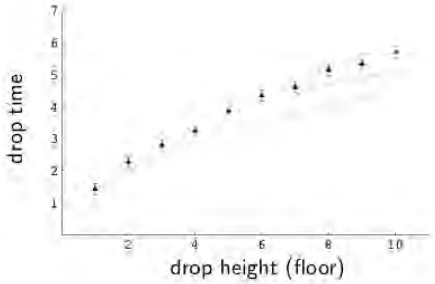

Accounting for modeling error:
This systematic discrepancy between the computational model and the experimental data is accounted for with a statistical adjustment. This term adjusts the model-based predictions to better match the data. The resulting 90% prediction intervals for floors 7 through 10 are also shown in this figure. Note the prediction intervals become wider as the drop level moves away from the floors with experimental data. The points corresponding to floors 7 through 10 show experimental observations taken later, solely for evaluation of the predictions.

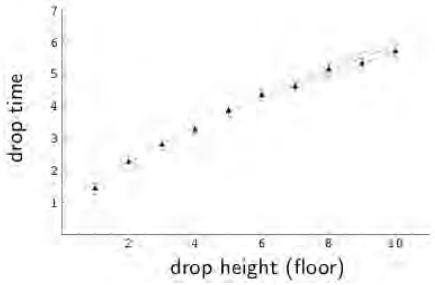

Improved mathematical model:
An improved simulation model was constructed which accounts for air resistance. A parameter controlling the strength of the resistance must be estimated from the data, resulting in some prediction uncertainty (90% prediction intervals are shown for floors 7 through 10). The improved model better captures the physics in the process, giving reduced prediction uncertainty.

Images courtesy of Dave Higdon (Los Alamos National Laboratory) and Richard Klein (University of California, Berkeley, and Lawrence Livermore National Laboratory).

However, interpreting these resulting probabilities can have its own set of difficulties. In practice, solving problems in inversion and computer model calibration can be complicated by a variety of issues, including high-dimensional parameter spaces, computationally demanding forward models, nonlinearity and/or complexity in the forward model, sparse physical observations, and inadequacies (numerical and physical) in the computational model.

Sensitivity Analysis

Sensitivity analysis is the systematic study of how model inputs—parameters, initial and boundary conditions—affect key model outputs. Depending on the application, one might use local derivatives or global descriptors such as Sobol's functional decomposition or variance decomposition. Also, the needs of the application may range from simple ranking of the importance of inputs to a response surface model that predicts the output given the input settings. Such sensitivity studies are complicated by a number of factors, including the dimensionality of the input space, the complexity of the computational model, limited forward model runs due to the computational demands of the model, the availability of adjoint solvers or derivative information, stochastic simulation output, and high-dimensional output. Challenges in sensitivity analysis include dealing with these factors while addressing the needs of the application.

Predictions from Multimodel Ensembles

In applications such as climate and weather prediction, a variety of different models are available for making a particular forecast. Predictions and estimates of uncertainty that combine the results of multiple models often give better predictions than any single model. Recent research has focused on approaches for combining results from different models. This includes approaches based on Bayesian model averaging as well as game theoretical paradigms. Weather forecasts based on such approaches have proven very successful (e.g., see the University of Washington's Probcast or the Canadian Weather Office's ensemble forecasts).

Representing Uncertainty

The question of how to represent and communicate uncertainties is a topic of research both from a practical and theoretical point of view. A fair bit of theoretical research is aimed at the mathematical calculus of uncertainty. This includes extensions and alternatives to standard probabilistic reasoning, such as Dempster-Schafer theory and imprecise probabilities. When uncertainties are needed for investigations requiring computational models, additional considerations arise. For example, if the simulation output is a daily surface-temperature field over the globe for the next 200 years, representing uncertainty and dependencies is complex. Should ensembles be used to represent plausible outcomes? How should these ensembles of simulation output be stored? How can high-consequence/low-probability outcomes be discovered in this massive output? Here some research investigations attempt to leverage theory that exploits high dimensionality to bound probabilities and system behavior. Finally, even when uncertainties are well captured, how best to communicate such uncertainties to the public or to decision-makers is also a topic of ongoing research.

Verification and Validation

Verification and validation has been a staple of the computational model assessment community for the past few decades. Standard definitions are below.

- **Verification** is the process of determining, as completely as possible, whether a computer code correctly implements the intended algorithms, and determining the accuracy with which the algorithms solve the intended equations.

- **Validation** is an assessment of the degree to which predictions of a code represent the intended physical phenomena, with the purpose of quantifying how accurately the model equations represent physical reality over a specified regime of applicability.

Verification and validation have focused on the comparisons: code versus math model and code versus reality. Clearly, both of these activities are related to UQ. Research in verification includes convergence assessment, estimation of bias and uncertainty due to numerical error, a posteriori error estimation, and the method of manufactured solutions. Validation focuses on comparing simulation-based predictions to experimental results. Research in validation includes choosing/designing experiments, assessing experimental uncertainties, propagating uncertainty, and determining the physical regime in which a code is validated.

Clearly, much of the activity in validation overlaps with UQ, and much of the research from this area is relevant to UQ. In many settings, verification and validation have a regulatory flavor. Hence, verification and validation may not focus on other questions that are clearly in the realm of UQ. For example, how can a model be adjusted to give improved predictive accuracy? How should different sources of uncertainty be combined to determine prediction uncertainty? How should one produce uncertainties for predictions outside of the validation regime?

Data Assimilation

A number of applications in monitoring and surveillance require persistent updating of the state of the system, predictions, and uncertainties based on continual or periodic collection of new physical observations. These data are combined with the computational model to update inferences. When the physical system is linear, and observation and model evolution errors are Gaussian, this updating can be accomplished with the Kalman filter, which updates the state of the physical system in an iterative fashion using the new data. More recent research has focused on updating large-scale, nonlinear systems—such as interacting particle systems, oceans, and atmospheres. Here the nonlinearities, massive data (e.g., satellite and sensor readings), and the extreme computational effort required to run these models have made it necessary to develop new approaches for updating information about the physical system and producing predictions with uncertainties. This has motivated much recent research effort on extensions of the Kalman filter, such as the ensemble Kalman filter, extended Kalman filter, and Monte Carlo-based techniques like the particle filter.

Optimization, Adaptive Design, and Feedback

These large-scale computational models are a powerful tool for planning and decision-making. Such models can be used to help assess important questions regarding the management of a particular system.

For example, how many and what types of sensors are required to monitor the planet's greenhouse gases over time? How should limited resources be used for this monitoring? Airborne measurements? Remote sensing? Ground-based sensors? Also, what are the likely responses to different mitigation actions? These management questions, which have analogs in almost any other application, require an understanding of uncertainties and of how various actions will reduce or otherwise affect these uncertainties. Research in this area includes optimization, decision theory, the design of experiments and optimal sensor networks, and optimal resource allocation. This line of research typically comes with a much larger computational burden because the impact of many possible strategies (e.g., mitigation, allocation, design, etc.) must be assessed to find one that is optimal (or near optimal). The difficulty of such problems is further complicated by the computational demands of the models, as well as by the model's potential to deviate from reality in ways relevant to addressing the optimization question.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

Role of Uncertainty Quantification in Extreme Computing

Extreme-scale systems present daunting challenges for their usage. Considering these challenges in the context of a UQ workload, researchers are presented with unique opportunities to pursue these challenges differently than with previous terascale and petascale systems and thus make these problems more amenable to solution.

The U.S. Department of Energy's (DOE) Exascale Initiative plans envision the development and deployment of exascale systems in the 2018 time frame. These systems will bring unprecedented computational power to pressing scientific simulation activities for both DOE's Office of Science and the National Nuclear Security Administration. However, the technological trends underpinning these systems will push the architectures in directions that will present many challenges to their gainful use on large-scale scientific discovery activities that are based on predictive scientific simulation.

Current computer industry trends portend significant challenges to scientific simulation in general. New programming models will probably be required in the exascale generation of platforms to deal with up to a billion-way parallelism and application resiliency challenges. Also, new usage model paradigms need to be considered. This is where UQ might make a very important contribution to researchers' usage model for exascale systems. For example, today's applications are considered as a stand-alone single job running across the entire system or space sharing the system with a few other large applications. This requires a very high mean time between failure for the hardware and software, because any nonredundant hardware failures cause the application to terminate abnormally. The most common way of dealing with the fact that the mean time between failure of today's systems (days to a week) is usually much shorter than the mean run-time of those applications (weeks to months) is for the application to periodically checkpoint the internal state of the calculation and restart from the most recent checkpoint upon abnormal termination. However, the cloud-computing model used for most Web 2.0 services (e.g., Amazon®, Google, Yahoo!®, eBay®[8]) uses a transaction model of service where each hypertext transfer protocol (HTTP) request can time out and be retried. If there is a hardware or software failure of a component causing a HTTP request to fail, another host can seamlessly take over servicing the subsequent requests.

[8] Amazon is a registered trademark of Amazon, Inc.
Yahoo! is a registered trademark of Yahoo! Inc.
eBay is a registered trademark of eBay, Inc.

That way, the overall service can have 99.999% availability but the mean time between failure of the underlying hardware can be fairly low. Indeed, Google is very proud of the "extremely inexpensive" hardware used for its cloud-computing infrastructure. Because future UQ "throughput" workloads on the order of 10,000 to 100,000 ensemble jobs may be typical, one can think of each job or instantiation of the predictive simulation application running on the exascale platform as a "transaction" that can be retried (or transparently restarted from a checkpoint) upon failure. This approach completely changes the application resiliency problem that must be solved for exascale systems.

Another challenge for ensemble-based UQ on exascale systems (and petascale systems before them) is the vast quantity of data that are generated during the ensemble runs. It is not just the amount of data (measured in exabytes) that is challenging, but the vast quantity of files and directories used to map data back to an individual run. Without a database tracking this mapping, it will be impossible to manage the ensemble data. Most modern codes are written in object-oriented languages (e.g., python, C++) in an object-oriented style. Most UQ frameworks, or "pipelines," are also written with object-oriented languages and techniques, but the object hierarchies are separate from the applications they manage. Also, several parallel file systems are object oriented and use object-storage devices (RAID [redundant array of inexpensive disks] devices or the disks themselves). However, the file system objects do not map onto the storage device objects. Thus, key information about the data to be stored (metadata) is lost.

One of many possible approaches to addressing these issues is to adopt a new file system paradigm that allows the UQ pipeline to define the objects that are used and augmented by the applications it drives. In addition, the packages within the file system could augment these objects and then pass them directly to the file system with the full context of the computation and input/output operation (e.g., on time step X for package A within application B working on parameter study Z for UQ study alpha) directly into the file system. Then the file system becomes an object-oriented database that allows one to search on application-defined metadata parameters consistent with the object-oriented hierarchy.

Additional challenges for UQ on exascale platforms, due to their vast scale, will be job management (e.g., job preparation, execution, updating, results analyses, and termination). The vast scale of exascale platforms will make this problem more difficult because of the quantity of jobs being run and the need to deal with a necessarily large number of problems and errors. A UQ capacity workload interspersed with "full system" UQ jobs will stress future job scheduling and allocations management infrastructure. However, the UQ capacity workload(s) could play an important "bottom feeder" role by absorbing otherwise idle cycles and by presenting a set of application transactions that can be terminated easily when "full system" runs need to be launched. Additional capacity jobs could then fill in the gaps caused by abnormal termination of the large job(s). In fact, if the scheduler could communicate with the UQ framework or pipeline and request jobs of various sizes with specific run times on demand—based on the current system load and expected future load, which reflects the queued jobs—then the scheduler could predictably fill empty run and/or time slots on the machine. This flexible run-time model could be used to keep exascale system use very near 100% capacity.

The above three examples show that considering exascale systems in the context of a UQ workload and in conjunction with a "full system" workload provides many opportunities for fundamentally changing the problems that need to be solved to make exascale systems deliver predictive simulation results with confidence. The demands of UQ on exascale machines lead to different and potentially simpler solutions than previous approaches for terascale and petascale systems with static large job workloads. Thus, UQ changes the approach for exascale systems and vice versa.

The above discussion focuses on the "UQ as throughput" system usage model. Here, a sequence of forward model runs—at various initial conditions and parameter settings—is used to understand sensitivities, propagate uncertainties, constrain parameter uncertainties, make predictions, and estimate probabilities. Figure 21 shows how a response surface can be used to interpolate model output and constrain parameter uncertainties. Depending on the problem, UQ as throughput will require from 10^3 to 10^8 computational model runs. These runs tend to be large in quantity and generate an exaflop/second computing load in aggregate. They also tend to generate vast quantities of data that need to be stored on a shared parallel file system and later analyzed as a single system or a tightly coupled simulation environment with a global parallel file system and low latency, high-bandwidth storage area networks.

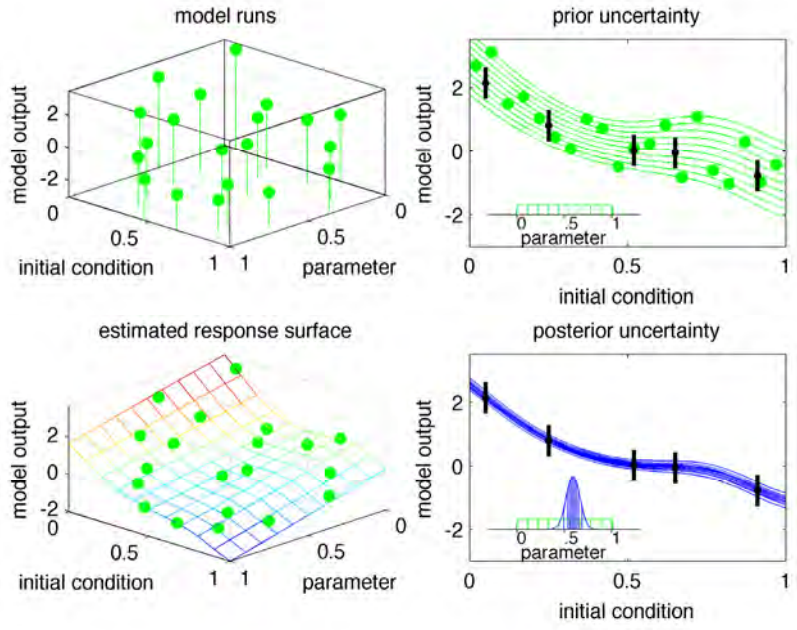

Figure 21. Parameter estimation and prediction uncertainty using response surface methodology with limited model runs. Top left: Model runs have been carried out according to a statistical design over the input settings determined by the initial condition and physical model parameter. The computational model output is given by the dots. Top right: The prior density for the unknown model parameter and the resulting simulations using the prior parameter uncertainties (green lines). The black dots show physical observations, and the corresponding black lines give 95% uncertainty bounds on the observations. Uncertainty quantification uses the data points and the model runs to reduce uncertainties in the parameter and in the prediction. Bottom left: A response surface estimate of the simulator output over the input space; the actual model runs are marked by the dots. Bottom right: The darker lines show the updated uncertainty for the model parameter and the resulting model-based predictions. Uncertainty in the model prediction is due to uncertainty regarding the parameter as well as uncertainty in the response surface estimate shown in the bottom left frame. Image courtesy of Dave Higdon (Los Alamos National Laboratory) and Richard Klein (University of California, Berkeley, and Lawrence Livermore National Laboratory).

On the other side of the spectrum is the challenge of conducting UQ on computational models that require extreme computing for just a single simulation. Here, one can expect only a handful of high-fidelity forward runs with which to carry out UQ. While this is impossible with most computational models, this may be possible if the next generation of computational models are constructed with UQ in mind. For example, computational models may be equipped with options for running faster, more-reduced models that produce more approximate results. Another possibility is to augment computational models with

adjoint solvers that can be used to compute derivatives of important outputs with respect to key inputs. In any case, for problems that require extreme-computing resources for a single run, UQ methods need to be co-designed with the computational model to allow exploration of sensitivities and uncertainties. Clearly, this is a research topic that should accompany the development of computational models for applications deserving of extreme-computational resources.

A Key Motivating Application for Uncertainty Quantification: Climate

An example of a scientific grand challenge problem requiring UQ analysis is the prediction of the future climate. The most important factor in climate prediction is Earth's equilibrium climate sensitivity (Bader et al. 2008). By definition, this is the increase in globally averaged surface temperature that would result if atmospheric carbon dioxide doubled; all other climate-forcing agents remained the same; and enough time elapsed for a new statistical steady state or "equilibrium" climate to be established. Equilibrium climate sensitivity may be considered a climate response per unit of climate forcing. It is important to practical applications of climate prediction because local climate impacts generally scale with globally averaged temperature change (Santer et al. 1990).

Svante August Arrhenius in 1897 concluded that Earth's equilibrium climate sensitivity is roughly 3°C. That estimate changed very little during the ensuing century. The National Research Council (1979) concluded that equilibrium climate sensitivity probably lies in the range of 1.5°C to 4.5°C, a statement repeated in subsequent Intergovernmental Panel on Climate Change assessment reports (e.g., McAvaney et al. 2001). In the late 1990s, work in the United Kingdom began that substantially advanced the state of the art of climate-UQ. This work employed large (~104) ensembles of climate model runs with differing input assumptions and produced, for the first time, probability density functions (PDFs) of equilibrium climate sensitivity (Meehl et al. 2007).

The PDFs are all consistent with the earlier, more qualitative estimates of equilibrium climate sensitivity cited above. Different and apparently equally sound methods, however, give significantly different PDFs. Furthermore, all PDFs produced to date are very broad—and as a result, the wide uncertainty range estimated over 30 years ago by the National Research Council has not been narrowed. In fact, higher sensitivity values (> 5°C) have non-negligible probability according to most of the PDFs. At the other end of the range, values < 1°C are implied by some calculations (e.g., Lindzen and Choi 2009).

Despite this rather unsatisfactory state of affairs, detailed probabilistic climate forecasts are starting to appear, beginning with the climate projections issued in 2009 by the United Kingdom Meteorological Office.[9] Clearly, a reasonable path forward must include sounder scientific underpinning for such climate predictions, which will be forced "by popular demand" whether or not scientists think they are ready to provide them. Therefore, researchers believe that UQ will play a high-priority role in climate science.

Comprehensive UQ studies in the climate domain require methodology to cope with high-dimensional uncertain input space and methodology to compare and contrast high-dimensional model output to observation-based data sets of various quality. UQ studies would yield understanding about the relationship between the uncertain climate processes, benefiting future climate-model development. UQ studies would also yield practical assessments of climate projection uncertainty, benefiting current climate-impact assessments. Propagating uncertainty in global climate model projections to regional

[9] See http://ukclimateprojections.defra.gov.uk.

scales introduces additional challenges because variations in regional fidelity (e.g., biases) between global climate model projections are all judged to be equal at the global scale. This applies particularly to regional precipitation. However, such regional UQ is necessary and crucial for practical climate-change impact studies and assessment of different climate-policy options.

PRIORITY RESEARCH DIRECTIONS

UQ is an emerging field that presents many research challenges. These challenges are best addressed by a focused effort to identify PRDs that researchers believe are opportunities to further advance the field. While UQ is a rather general topic with many possible research directions, the following comments focus on research directions that are of particular importance in light of extreme-scale computing.

Foundations in Uncertainty Quantification

Everyone has an intuitive notion of doubt or a lack of certainty regarding a possible outcome of an anticipated event. This is particularly important to decision-makers who, using national security applications, must assess the risks associated with numerous adverse events and take action with limited resources. For repeatable events, uncertainty statements can be assessed by comparison of actual outcomes. For well-calibrated UQ methods, 90% intervals should contain the actual temperature 90% of the time, in the long run. UQ methods that make better use of available information and computational models can produce much tighter intervals that still cover 90% of the time (Figure 22).

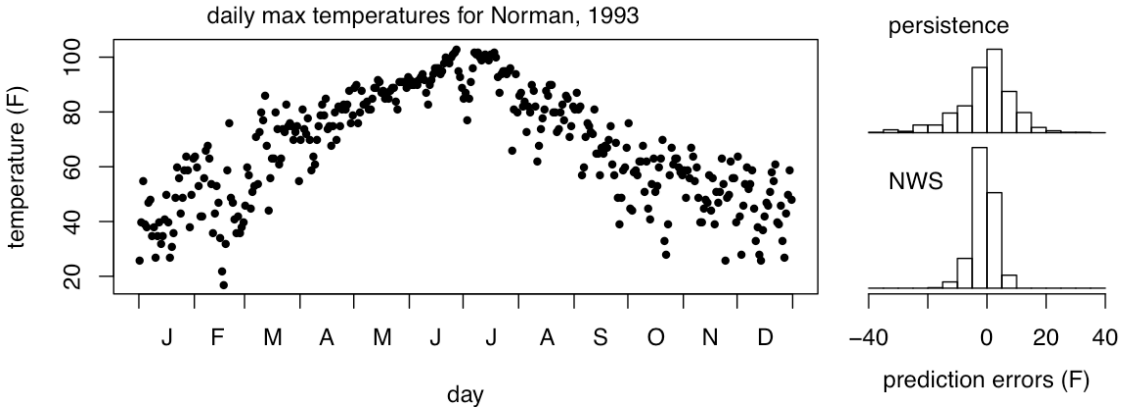

Figure 22. Daily maximum temperatures for Norman, Oklahoma (left), and histograms of next day prediction errors (right) using two forecasting models: persistence—predict tomorrow's temperature with today's temperature; and the National Weather Service (NWS) forecast. 90% of the actual temperatures are within ±14°F for the persistence forecasts, and ±6°F for the NWS forecasts. The greater accuracy of the NWS forecasts is due to its use of computational models and additional meteorological information. The assessment of these two forecast methods is straight forward because of the large amount of replication. These data are from Brooks and Doswell (1996). Image courtesy of Dave Higdon (Los Alamos National Laboratory).

While this type of assessment is quite intuitive and sensible, assessing predictions and uncertainties for rare, or "one of a kind" events is much more problematic. For example, what is the chance of a substantial terrorist attack in a major U.S. city? What is the chance the ocean level rises by 3 m within the next 40 years? What is the chance of a sizable asteroid impacting the Earth in the next 1000 years? Assessing such uncertainties is difficult; nonetheless, decision-makers still must decide what resources

and actions will be devoted to mitigating these types of risks. An assessment of the uncertainties associated with such events is fundamental to these decisions.

Scientific and Computational Challenges

Scientists expect that large-scale computational models, which encode numerous physical laws and phenomena, are viable tools to help with this uncertainty assessment. These models are clearly more than empirical, statistical formulations. There is a need for strong foundational underpinnings to make the best use of these computational models for important decisions regarding difficult-to-assess events. This PRD focuses on the development of language and framework for linking these promising computational models to reality, leading to meaningful statements of uncertainty.

The basic notion that better models give better predictions has not been formalized or mathematically encoded for assessing uncertainties in predictions from computational models. Current approaches, such as those in statistics and machine learning, make no theoretical distinction between a physically motivated computational model and an empirical, descriptive model. The better model eventually proves itself on physical observations that have been held back to test the competing predictions. The challenge is to develop a framework that better accounts for the nature of the model being used to make predictions, leading to meaningful prediction uncertainties. This framework should be helpful even when the physical data available to train and assess these models are limited.

Summary of Research Direction

The issues with UQ are present at all scales of scientific computation, from the desktop to the exascale, and include both descriptive and quantitative elements. A consistent treatment of uncertainty requires a consistent semantic basis, a language by which concepts regarding uncertainties in the context of physically based computational models can be communicated.

In addition to the semantic and mathematical fundamentals, a disciplined process for performing, documenting, and assessing UQ is required to support decision-makers. This process must include the universal glossary of terminology and standard guidelines for analysis and documentation, including descriptions of the problem, assumptions made, and methods applied. Efficient application of a standardized process is facilitated by examples of its application to relevant problems.

Expected Computational and Scientific Outcomes

The products of this research direction are frameworks, theories, and methods needed to perform UQ coherently and make the results meaningful and useful to decision-makers. These products include the following:

- methods for assessing the usefulness and quality of UQ analyses
- methods for analyzing the sensitivity of UQ to the assumptions and methodology
- methods for assessing uncertainties of model-based predictions in new, untested regimes (i.e., “extrapolations”)

- methods that leverage theoretical considerations and computational models to assess the likelihood of extreme, high-consequence events
- reporting guidelines for estimates and uncertainties, including disclosure of assumptions and methods
- compelling examples of UQ done well in problems with different degrees of complexity.

Potential Impact on National Security

Achievement of the research goals will result in consistent, quantified support for decision-makers. This research will also lead to uncertainty assessments with higher quality and more common features. The variety of applications of such support covers the entire spectrum of activities supported by simulation at any scale up to and including the exascale. This research should help illuminate when high fidelity, first-principles models are required and when they are not. Examples of potential applications include stockpile stewardship, nuclear reactor safety, and many other aspects of the nuclear security enterprise; considerations of legislative and judicial options for actions to address climate change; and the use of very large-scale computing as a surrogate for experiments of discovery in basic science.

Countering the Curse of Dimensionality

Scientific and Computational Challenges

The current state of the art for UQ science of multiphysics simulation codes is performed with an ensemble-of-models approach. That is, researchers take a model of interest and identify the subset of input quantities (usually ~ 7 to 10) they estimate will dominate the uncertainties in the predictions of interest because researchers have limited computational power. They then compute the model thousands of times with differing combinations of input quantities (ensembles) and constrain some or all of the inputs by using available data with their own associated uncertainties. Finally, researchers compute bounding uncertainties for predicted model outputs. Unfortunately, there are many limitations to researchers' capabilities—and in particular, one stands out. To provide more accurate uncertainty bounds, researchers need to be able to vary all of the input quantities with uncertainties that can influence a given prediction. Generally, researchers know that of an order of 100 input parameters (e.g., climate, inertial confinement fusion) in many multiphysics areas can influence simulation predictions. Thus, to advance UQ science, researchers need to develop new methodologies that can accommodate the vast number of existing uncertainties (dimensions) with a minimum of code calculations. This challenge is referred to as the "curse of dimensionality." To do so in an efficient manner, researchers need an automated process (UQ pipeline) that can self-guide the large numbers of ensemble simulations needed to explore the complexity of uncertainties in many tens to hundreds of input parameters, and that can identify and apply the most efficient combination of methodologies for minimizing the number of code runs required. Such a UQ pipeline will enable researchers to test and refine the UQ methodologies as they are developed, and it will eventually enable efficient use of exascale computing. Current practice to reduce the numbers of simulations required is to constrain parameter ranges (i.e., input space) based on available observational data, physical considerations, and/or the results of previous studies. Two simple and commonly used approaches to constructing the ensembles are Monte Carlo and Latin hypercube sampling of the input space. Response surfaces (also known as a statistical response model, a surrogate model, or a meta model) are then constructed from the ensemble results, and these are then convolved with observational data to further constrain input parameters and to create uncertainty bounds on model

outputs. This approach works well for studies limited to the variation of a handful of input parameters. However, this approach does not hold when the simulator is computationally expensive and when the uncertain input-parameter space is high dimensional. This is the case for comprehensive UQ for multiphysics, multiscale codes (e.g., global climate models).

Because the problems scientists are interested in have uncertainties associated with tens to hundreds of input parameters, they have the problem of high dimensionality. For computationally expensive simulators and current sampling methods, this makes comprehensive UQ intractable with current and even future envisioned computational platforms. Put simply, Monte Carlo and Latin hypercube sampling techniques cannot adequately resolve the response surface when the input space is composed of tens of dimensions. For high-dimensional input spaces, it is nearly impossible to design a computer experiment up front that captures all important aspects of the input-output relationship, which has motivated various adaptive (sequential) sampling strategies.

Summary of Research Direction

There is a need to reduce the effective size of the input space, either through formal dimension reduction techniques or through input variable selection methods. Promising research directions are given below.

- **Self-adaptive exploration of response surfaces.** Methods to produce an efficient and robust ensemble of simulation results through adaptive sample refinement (ASR) need to be developed. These methods may guide the ASR process through topological characterization and also through deficiencies in the predictive accuracy of the response model. Level sets provide isoparametric contours of the response function. When the model is high dimensional, finding the boundary of such sets accurately is extremely difficult. Methods need to be developed to refine adaptively near the boundary of the set.

- **Large-scale parallel response surface analysis methods.** Global response surface approaches to UQ that use an ensemble-of-models approach need to be broadened. Research is required on the generalization and extension of regression models to include basis functions with local support for the purpose of enhancing an ASR capability. New parallel, scalable algorithms for the purpose of generating response functions of very high dimensionality are also needed.

- **Theory/methods for high-dimensional representations.** Contour tree approaches can be developed to analyze the topology of high-dimensional functions. These can be used to improve ASR by identifying monotone/nonmonotone regions to guide sampling of high-dimensional response surfaces. New approaches would include basis function transformations. Specifically, the objective is to transform the original high-dimensional parameter space into a reduced or lower-order space that can be used as a surrogate for the original input space to reveal underlying structure in the original space.

- **Surfing the UQ pipeline.** Future exascale studies will likely consist of tens of thousands or greater numbers of ensemble simulations, and the sample space studied will be characterized by complex, high-dimensional local and global structure. This class of UQ studies will be intractable if the user is required to make decisions concerning the guidance of the UQ study. This complexity requires a UQ pipeline to include self-guiding, self-adapting technologies that steer the ensemble of simulations, without user interaction, toward the areas in the sample space where the effort should be focused.

Expected Computational and Scientific Outcomes

With the development of advanced methodologies to address the problem of dimensionality, it will become possible to perform UQ analysis across a broad range of multiphysics, multiscale scientific problems that include the main uncertainties inherent in the underlying physics models, numerical algorithms, databases, and inputs and output observables. Such developments will also result in a self-adapting, self-guiding UQ pipeline that will enable UQ studies to be performed on exascale platforms.

Potential Impact on National Security

Progress in key areas of UQ research—such as the curse of dimensionality—will affect critical areas of importance to national security such as nuclear weapons and stockpile stewardship science, climate prediction, and inertial confinement fusion to name a few. In the area of climate prediction, it will likely be possible to make consistent uncertainty estimates in global climate sensitivity, predict regional climate impacts, and move to exascale computing within 8 years to include vastly improved cloud physics.

Intrusive/Embedded Uncertainty Quantification

Scientific and Computational Challenges

There are three fundamental components to end-to-end UQ for large-scale simulations, whether in the form of partial differential equations (PDEs), ordinary differential equations (ODEs), integral equations, discrete particle systems, or other simulation models. These components include the following:

- Statistical inverse problem: the estimation of uncertainty in model parameters or model structure from observations or measurements.

- Uncertainty propagation problem: the propagation of input-parameter uncertainties through the simulation model to predict model outputs.

- Stochastic optimization problem: the solution of optimal design or control problems that are governed by the stochastic forward problem and make use of statistics of model predictions as objectives and/or constraints.

Many contemporary techniques for solving stochastic inverse, forward, and optimization problems suffer from the curse of dimensionality and become computationally intractable for problems governed by large-scale simulation models with high-dimensional uncertainties. The availability of exascale computing by itself will not overcome these challenges; researchers need fundamentally new algorithms and analysis for estimation, propagation, and optimization under the presence of uncertainties in large-scale simulations of complex systems.

Summary of Research Direction

It is believed that one of the keys to overcoming the twin curses of high dimensionality and expensive forward simulations in UQ methods is to exploit the structure of the mathematical model that maps parameter inputs to output quantities of interest. Most contemporary UQ methods, such as conventional Monte Carlo methods, treat this input/output map (i.e., response surface) as a black box. Yet recently

developed methods that exploit this input/output map structure have been critical to the solution of deterministic inverse and other optimization problems with millions of parameters at a cost of a handful of forward simulations. Recently developed methods are also critical to the large reduction of the cost of sampling by approximating response surfaces and constructing reduced-order models as surrogates for expensive forward simulations. More recently, similar ideas to exploit the mathematical structure of the input/output map have begun to appear in UQ methods for forward and inverse propagation of uncertainty. Many of these methods employ rapidly computed derivative information, motivated by the fact that for most systems governed by differential (and related) equations, the outputs are locally smooth and thus derivative information is generally useful. Researchers refer to such methods as "intrusive" or "embedded" methods because they require access to and analysis of at least the Jacobians of the underlying forward operators. Several examples of new research directions in UQ based on intrusive methods are described in the following paragraphs.

Langevin methods for sampling probability densities, whose trajectories are driven by gradients of the (log of the) target density, are beginning to be employed for high-dimensional PDE-based inverse problems. Computation of the gradient is greatly facilitated by the use of adjoint equations (whose operator is the adjoint, or transpose, or the linearized forward operator). The gradient can be computed at a cost of at most a single-forward simulation, and usually less, because the adjoint equation is always linear even when the forward problem is nonlinear. The cost of this capability is that legacy-forward simulation codes often have not been designed to compute adjoints, and retrofitting complex legacy codes with adjoint capabilities entails a significant refactoring. Hessians (i.e., second-order sensitivities) provide even richer information than gradients and play a critical role in identifying significant directions in high-dimensional inverse problems (i.e., those directions in parameter space for which the data provide meaningful information about the model). Moreover, recent techniques for building low-rank approximations of data-misfit Hessians based on analysis of the underlying infinite-dimensional operators have allowed for significant acceleration of sampling methods in inverse problems.

Hessians are also beginning to be used in the construction of reduced-order models as surrogates for expensive-forward simulations; Hessians can efficiently steer the placement of design points in parameter space at which full-order model outputs are computed and employed for reduced model construction. Low-dimensional reduced models constructed in this way have proven to be very effective at approximating full-order model outputs; in turn, they have facilitated rapid sampling of probability densities that embed expensive forward simulation. Similarly, Hessians have begun to be used in the construction of Gaussian process approximation of response surfaces, both in effectively placing design points in parameter space as well as in informing the Gaussian process approximation.

Polynomial chaos methods are another technique for exploiting the mathematical structure of the parameter-to-output map; in this case, by approximation by multivariate polynomials. While convergence can be very fast, standard polynomial chaos methods suffer from the curse of dimensionality and they soon become unaffordable. Sparse-grid approximation techniques improve the convergence of polynomial chaos methods, and even greater improvement can be obtained using anisotropic sparse approximation, which exploits the relative importance of different input random variables on the solution.

Still, despite very promising performance exhibited by emerging intrusive UQ methods for some challenging problems, it is fair to judge such methods are in their infancy, and substantial works lies ahead in extending, improving the robustness of, tailoring, and scaling up these methods to address the complex, large-scale, nonlinear multiphysics, multiscale, high-parameter-dimension UQ problems arising

in national security applications. To do this, fundamentally new ideas are needed to further exploit mathematically and computationally the functional relationship between input parameters and output quantities of interest.

Expected Computational and Scientific Outcomes

As discussed above, current methods for quantifying uncertainties in simulation models are incapable of scaling up to expensive simulations characterized by large numbers of uncertain parameters. The development of UQ algorithms and methods that can exploit input/output structure to overcome these pervasive and abiding barriers will enable computational scientists to conduct large-scale simulations with quantified uncertainties, thereby transforming the ability to effect meaningful predictions for many critical scientific, societal, and strategic problems.

Potential Impact on National Security

One of the central challenges facing the field of national security is to employ large-scale simulation as a tool for decision-making involving uncertain complex systems. For such problems, the "single point" deterministic predictions produced by contemporary large-scale simulations are of little use for decision-making; to be useful, these predictions must be accompanied by estimates of their uncertainty. Many problems in the national security portfolio are characterized by large-scale, expensive simulations and high-dimensional parameter spaces. Intrusive UQ methods promise significant breakthroughs in researchers' ability to address high-dimensional uncertainty and expensive simulation models. By permitting predictive simulations to be employed in a tightly integrated way, success in developing such techniques will lead ultimately to a revolution in the way that decision-making under uncertainty is conducted.

Uncertainty Quantification in Data-Rich Environments

The role of data in terms of size, quality, and access is becoming ever more prominent in all areas of human activity. Dramatic changes underlying the data revolution have been achieved by rapid progress in solid-state technology, ubiquitous networking and sensing, cheap storage, and associated advances in the computer and information sciences. While no one doubts the many uses of the analysis of data sets that previously could not be imagined, one must handle the "data flood"—a runaway generation of data that can only be contained by a corresponding exponential increase in researchers' abilities to ingest, store, organize, and interrogate the data stream. From the analysis perspective, it is useful to think about two extreme classes of problems: those related to real-time (or near real-time) applications; and those related to dealing with very large data sets, both observed and simulated. In both cases, UQ must play an essential role in confronting the data flood.

Scientific and Computational Challenges

A significant number of scientific and computational challenges are posed by large data sets and data throughput. Areas in which these arise include astrophysics, biology, climate modeling, cyber security, Earth sciences, nuclear and particle physics, and situational awareness. Petabyte databases and data gathering rates of tens of terabytes/day are already with us; a thousand-fold increase may be envisioned well within the next decade. Challenges that are particularly relevant to UQ include the following:

- **UQ for Approximate Algorithms**. One of the key challenges is developing very fast algorithms for data analysis, whether statistical in nature or related to aspects such as graph and network analysis and pattern recognition. The severity of this problem is such that even O(NlogN) algorithms may be far too slow, and O(N) or even O(1) algorithms may be needed. However, the only available fast algorithms are likely to be approximate ones in which case building a robust UQ infrastructure for the predictions from these algorithms is essential.

- **UQ and Extreme Modeling and Simulation I.** The data sets produced by extreme-scale computing will soon be as rich as experimental or observational databases. Indeed, large future experiments and observational campaigns are already being designed in a feedback loop with simulations. This imposes a serious requirement on the validity of the modeling and simulation process that in turn requires a new class of UQ methods able to deal with problems such as high dimensionality and predictions for extreme values.

- **UQ and Extreme Modeling and Simulation II.** Very fine-grained and complex simulations have a natural application as test beds for problems that cannot be handled experimentally (e.g., disaster response). In cases such as this, it is unlikely that any single computational model can be realistically accurate. However, the fact many intervention strategies (and many models) can be separately investigated immediately poses the need for a UQ paradigm able to assess the usefulness and risks of these strategies, given the underlying limits of the models and simulations.

- **UQ Co-design with New Computing Architectures.** It is widely accepted that supercomputing architectures will undergo major changes over the next decade. Next-generation applications must deal with severe concurrency and latency challenges and must be developed in close concert with the evolving architectures. This will also certainly be true of UQ methods and frameworks, which must be flexible enough to encompass next-generation hybrid supercomputers, data-intensive supercomputers, very large cloud-computing platforms, and special-purpose machines.

Summary of Research Direction

This research direction is driven by the proliferation of data, both sensor-derived and digitally archived, as well as generated by large-scale computers. In many cases, the scale of the data throughput and size renders answering classes of precise questions a meaningless exercise, even if the data are well characterized. Thus, one task imposed by this new arena is to develop UQ strategies for inherently approximate analyses. In addition, the scale of the data is such that it will contain (possibly) very high-dimensional dependencies, and dealing with them will require robust UQ-controlled techniques of model and data reduction, including controlling uncertainties from combining large-volume disparate data sources. Finally, the UQ methodology will be required to be sensitive to extreme-scale computer architecture (both conventional and data intensive) because it will evolve considerably.

Expected Computational and Scientific Outcomes

This research direction will be an essential aspect of information extraction from large data sets, especially as these data sets quickly scale beyond the reach of current analysis methodology. Many of the applications will be associated with major scientific and engineering efforts, national security data gathering and databases, and important social issues—all areas where a reliable and robust UQ paradigm should be considered a key, if not the dominant requirement.

Potential Impact on National Security

The impact on national security is very significant. Decision-makers will need to understand the ramifications of various actions—undertaken both by them and others—in an increasingly complex and data-rich environment. The ability to extract information, features, and dependencies, and to be able to quantify the uncertainties associated with certain actions will be a key aspect of UQ for national security. In addition, at a lower level, many aspects of national security will require dealing directly with UQ issues for data sets at extreme scales; cyber security, counterterrorism, disaster response, and situational awareness are obvious examples.

Combining Disparate Models and Data Sources

With many physical systems, there is no integrated computational model or code that incorporates all relevant processes from which scientific inferences can be made. Typically, a number of computational models are available to model different aspects of the system. Although they may share some commonality, these separate models, typically focus on different aspects of the system. Similarly, wide varieties of data sources are often available to inform about a given physical system. A PRD is the development of a conceptual framework and methodology for making scientific inferences with the aid of these disparate models and data sources.

Scientific and Computational Challenges

There are a number of scientific investigations of importance to national security that motivate the need to combine disparate models and data sources. These include the following:

- **Inferring material behavior.** Computational models for materials are now available on multiple scales of resolution, from atomic to meso- to macroscales. The atomistic models, although computationally demanding, are very nearly first-principles models. Hence, they can be used to help infer bulk properties of the material, such as equation of state or strength. The development of UQ methods are essential for combining these models to infer bulk material properties at temperature and pressure conditions that cannot be accessed in laboratory experiments, particularly for furnishing realistic model uncertainty estimates appropriate to such extrapolations.

- **Estimating, tracking, and managing greenhouse gas (GHG) fluxes.** While no fully integrated Earth-system model exists (or is likely to exist) that incorporates all processes relevant to atmospheric GHGs, there is a wide variety of models and data sources to aid in the estimation, tracking, and management of GHG fluxes. Relevant models might include atmospheric transport, ecological dynamics, the carbon cycle, land use, social behavior, and energy infrastructure. Relevant data sources are equally varied—econometric inventories and summaries; census information; land, sea, and air-based sensors; satellite observations; and isotopics are just some of the data sources. A framework and methodology is clearly needed to infer the current state of GHG fluxes. These are also needed to infer the impacts resulting from potential mitigation strategies.

- **Inferring climate change at the regional level.** While climate change predictions are typically made with large-scale global circulation models, planners need to know how such changes will affect their local climates. For example, will current water sources dry up? Will agriculture remain tenable? Approaches exist for downscaling coarse-level climate change information to the regional

level, but UQ for these regional predictions is not well developed. UQ methods for this application require the combination of global and regional models with local information such as weather, topography, groundcover, and hydrology to predict the effect of climate change at this local level.

In addition to the specific above-mentioned application areas, relevant scientific challenges exist in nonproliferation where a wide variety of information is available (from intelligence to sensor signals), and models may include agent-based or sociotechnical simulations.

Summary of Research Direction

Just as research in "data fusion" focuses on the development of new frameworks, methods, and algorithms for making use of diverse information sources for making inferences, this research direction extends this notion to incorporate the use of diverse computational models as well. While the construction of tightly coupled multiphysics codes tackles this problem head on by fusing multiple models into a single code, this solution is labor intensive and very specific to a particular application area. Also, management and upkeep for such codes can be a daunting task. This research direction proposes the development of new UQ methods to link different models and data sources into a common inferential framework.

This looser coupling of different models will necessarily lead to a greater reliance on physical observations. Hence, the ability to incorporate all available data and information sources—possibly in a dynamically updating situation—will be crucial. These sources may include physical observations, experiments, and expert judgment. New statistical approaches for combining information are expected to be developed to make the required inferences. On the computational side, such inferences will undoubtedly require many model runs from each of the models involved with the analysis. It is also expected that approaches will be developed to determine what combination of model runs are required and how to optimize their allocation to high-performance computing resources.

Expected Computational and Scientific Outcomes

If successful, this line of research will greatly improve the ability of different data sources and computational models to make inferences about complicated physical or social systems. Such research will also give insight on how to improve the resulting inferences. For example, which models should be improved or altered? What will be an impact of including additional data sources or computational models?

Potential Impact on National Security

The result of such research will be an agile framework for constructing available and relevant models and data sources to answer important questions relating to national security. The ability to assess system response with uncertainties—based on the available models and data—will be crucial for decision-makers to formulate responses to potential threats to national security.

Uncertainty Quantification for Emergency Response

Future national security threats include events precipitated by weapons of mass destruction, natural disasters (including weather-related calamities), manmade disasters, pandemic events that would have biological impacts, and potential threats to the population caused by human activities that could result in

significant changes in the ecosystem. Effective, timely response to such threats requires accurate assessment of the scenario and its possible evolutionary state, coupled with useful estimates of the risks and consequences associated with potential response options. Accurate assessment of a given threat scenario or a particular response option can be an extreme computational challenge by itself. However, to be of the most use to emergency responders, such an assessment must be timely and must include an estimate of the uncertainties associated with the various possible evolutionary states and potential response options. Therefore, quantifying uncertainties for emergency response adds another dimension to the requirements for exascale computing.

Scientific and Computational Challenges

Predictive modeling of disaster scenarios and response options requires advancements in physics models and computational capabilities. The challenges associated with the physics modeling capabilities are similar to those discussed elsewhere in this panel report. The challenges associated with the UQ aspects of emergency response are related to the near real-time decision-support requirement for effective response. One challenge is the integration of the scientific modeling capability with historical and prompt data describing the scenario. Such data include the following:

- source term (nuclear, chemical, biological, and environmental)
- emplacement environment (buildings, structures, etc.)
- geographical terrain and features
- local and regional weather conditions at the time of release
- population density and distribution
- available response and mitigation options.

The challenge for UQ is to estimate the variability in the consequences of the response options from the potential variability of model predictions based on the uncertainties in the scenario information. The large number of variables associated with each of the inputs, as well as the coupling between the associated models needed to determine consequences in real time, requires a very large number of calculations at large scale.

Summary of Research Direction

There is a spectrum of possible approaches for developing the capability for emergency-response decision support based on physical modeling with quantified uncertainties. The spectrum runs from the exclusive use of precomputed databases of scenario and response options and the associated uncertainties to near-real-time modeling and UQ on demand. The precomputed database end of the spectrum suffers from the obvious drawback of not accurately reflecting a particular threat and currently available response options, thereby adding uncertainty to those previously estimated. The near–real-time modeling approach requires more computing capability than will be practically available—even at the exascale. Therefore, a practical capability will require some combination of the two approaches. The proposed approach is to couple databases of historical data and high fidelity, precomputed scenarios, response options, and the associated

uncertainties with current information to produce accurate, scenario-specific support for emergency response decisions.

Therefore, the focus will be on the codevelopment of precomputing strategies with novel analysis techniques that can efficiently exploit existing information databases (e.g., maps, terrain, buildings, etc.), prompt data (such as weather conditions and availability of emergency services), and computational model runs to give critical information regarding the scene and to evaluate and propose potential responses. Particular activities might include the following:

- establishing detailed worldwide databases of terrain, buildings, and population distribution for baseline calculations based on satellite data mining
- developing a capability to couple weather or seismic data in real time to address structural impacts and/or dispersion events anywhere in the world
- creating a comprehensive database of source-term calculations under a variety of emplacement geometries to predict patterns of dispersal and damage for weapons of mass destruction
- establishing a surrogate modeling capability for interpolation between precomputed, high-fidelity physics models and propagation of uncertainties
- developing methods that use simulation models of mass-evacuation strategies to optimize moving personnel to safety in a crisis
- developing approaches to estimate structural response for a variety of classes of buildings and using satellite data to populate those classes.

Expected Computational and Scientific Outcomes

The primary product of this PRD will be the development of capabilities necessary to support an accurate, timely emergency-response decision-support system. Such capabilities include those listed above and might include additional products such as the following:

- development of satellite photo data-mining techniques to construct local structural, traffic, and population models of target areas
- development of a seamless ability to couple real-time weather to the geographic database to calculate dispersion and effects.

These capabilities could be integrated into a nationwide decision-support system that combines laptop-based software for emergency responders coupled with remote database capabilities and computational resources.

Potential Impact on National Security

Risk-based emergency-response decision support will enable emergency responders to optimize resources and minimize consequences of natural and manmade disasters. This capability also provides accurate information to the national security infrastructure for assessments of impacts on national security and appropriate marshalling of national security assets.

CROSSCUTTING CHALLENGES

Given the very broad scope of this workshop, many common intentions and needs emerged, both in terms of scientific frontiers and methodologies, as well as computational resources and enabling research and development infrastructure. The principal crosscutting findings are as follows:

- The complexity of the simulations enabled by extreme-scale systems poses the need to develop methodologies for addressing new scales of interdisciplinary complex systems (in scientific disciplines and in urgent missions such as coupled energy-climate infrastructure, nonproliferation, etc.) to make improved prediction, control, and design tools available to decision-makers.

- Extreme-scale computing will continue to need to serve both huge computations at science frontiers (higher resolution, three-dimensional), and methods and models for enhanced predictive accuracy of both component processes and integral systems.

- Developing methodologies and algorithms capable of bridging length and time scales that span many orders of magnitude is an essential requirement for many applications anticipated at the exascale.

- It is important for all applications to advance and use rigorous uncertainty quantification, error analysis, optimization, inverse tools, etc. For example, this is critical for better quantification of metrics (performance, reliability, safety margins, yield) for "predictive" system modeling and simulation.

- A clear need exists to further develop and incorporate information science and technology methods to extract insights from the voluminous output of simulations and observational data; e.g., image analysis, machine learning, intelligent coarse-graining, stochastic modeling, model/hypothesis discovery and training.

- Computational algorithms/codes for relevant physical phenomena are needed, optimized to extreme-scale computing architectures:

    – robust, portable, modular, scalable code approaches

    – efficient connectivity among code modules

    – discrete and continuous phenomena, data, and models

    – combining theory- and data-driven modeling and simulation

    – data management tools for massive (heterogeneous) data (e.g., are there practical middle ways between compute- and data-intensive high-performance computing?)

    – determining and using accurate interactions for extreme/reactive environments

    – field-deployable supercomputing

    – hybrid simulation approaches (e.g., atomistic/continuum, quantum/classical, short/long range interactions)

  - covering vast space and time ranges for the multiscale and multiphysics realities of complex systems (far from equilibrium, nonlinear, nonadiabatic, metastable, heterogeneous, stochastic, rare-events, anomalies, long-tail distributions, defects, interfaces, hotspots).

- It is important to create and enhance mechanisms for deliberate teaming of expertise (application domains, computer science, applied math, etc.) to accelerate agile algorithm and code design for future extreme-scale computing architectures. This teaming must also be effectively coupled to a deliberate "co-design" strategy with major experimental and data facilities.

- Visionary partnerships (between universities, DOE laboratories, and industry) will be essential to achieving the envisioned extreme-scale computing impacts on national security science and mission frontiers, as well as to training a future workforce with the necessary interdisciplinary skills.

- Several areas of common interest to DOE's Office of Science and NNSA (e.g., extreme computing, climate-energy systems, materials under extreme conditions) were identified where progress should have major spin-off benefits to the U.S. industry.

CONCLUSIONS AND RECOMMENDATIONS

CONCLUSIONS

Extreme-scale computing will enable the study of more complex phenomena and increase the fidelity and accuracy of predictions. Furthermore, these improved simulations will be obtained in less elapsed time, which is important for many national security applications. Even on existing computer architectures, simulations of the more complex problems would require new mathematical models, numerical algorithms, and advances in technologies for uncertainty quantification and verification and validation. In addition, the architecture and software of the first generation of extreme-scale systems will pose major challenges and require new algorithms, programming models and languages, and system software features and capabilities. The latter challenges will be ameliorated if they are addressed with an integrative co-design approach that starts by obtaining a deep understanding of the scientific challenges that is used to guide (to the extent possible) the design of the extreme-scale hardware. Fortunately, many assets are available to meet these challenges—particularly research staff with highly relevant expertise, and teams that have embraced multidisciplinary approaches for advancing computational science, engineering and technologies, due in large part to the National Nuclear Security Administration's Advanced Simulation and Computing (ASC) Program and the U.S. Department of Energy's (DOE) Scientific Discovery through Advanced Computing (SciDAC) programs.

High-performance computing is entering a period of generational change that poses great challenges (even technologically) to the computing industry and to scientific users but also tremendous opportunity to raise the co-design of national science and technology (S&T) assets to a qualitatively new level of impact. Seizing this opportunity will allow the "scientific method" to meet the great national and societal systems challenges we face now and for the foreseeable decades. President Obama and his administration are investing resources and confidence in the scientific community to meet those challenges with innovation and science, technology, and engineering (ST&E) breakthroughs—breakthroughs of a significance that are likely to change scientific frontiers while addressing the great national security frontiers. This provides the ultimate timeliness requirement for investing in extreme computing, new experimental and data-gathering facilities and capabilities, and enabling the emergence of a new generation of interdisciplinary scientists, to reset integration and collaborative frameworks to meet national challenges—starting now.

RECOMMENDATIONS

A number of recommendations emerged from each of the panels; these were combined where possible in the following general recommendations.

Co-Design

Extreme-scale computing will bring many of the disciplines and systems frontiers discussed in this report to an era of greatly enhanced discovery, predictive, and design capability. The scientific community's ability to predict outcomes with far greater fidelity and with quantified uncertainties will be key metrics of success. This exciting goal can only be achieved by a deliberate strategy of advancing key enabling technologies in theory and modeling, in experiments and observational tools, and in simulation science. However, it also depends critically on similarly deliberate strategies of "co-design" of national assets—

including designing experiments and measurements for collection of relevant data— and collaborative infrastructures to use them. This will accelerate feedback and time to result, but more importantly change the scale of challenges to which the scientific method can aspire.

Algorithms and codes need to be agile and component-structured for the single-physics and integrated packages for the ST&E problems envisioned in this report. Next-generation architectures must be able to handle these codes. Vice versa, the algorithms and codes must take advantage of the architectures that new technology can permit. Early joint planning is essential. There is perhaps an even greater generational opportunity to align this co-design with planning for the next generation of experimental/data facilities, so that the data being acquired have maximum relevance to the theory- and data-driven models that extreme computing is simulating. Again vice versa, theory, modeling, and extreme computing must have the greatest synergy with what can be measured because of new experimental and engineering technology.

Development of, Deployment of, and Access to Extreme-Scale Computational Resources

As detailed in the panel reports, the discipline and mission frontiers described require extreme-scale computing. Extreme-scale computing facilities should be designed to accommodate the needs of several and diverse communities. Associated with these facilities, there must be a balanced and complete infrastructure to make them accessible and efficient, and primed for the decadal challenges—including applied mathematics, computer science, algorithms, optimization, embedded uncertainty quantification (UQ) capability; and supporting software, data storage, data analysis, visualization, and associated analytical tools. Early access—and indeed the consequent co-design opportunities—for application and domain experts is essential if next-generation architectures are to be optimized for next-generation applications.

The panel topics covered in this report clearly make the case for different computational platforms. It is unlikely to be efficient or even feasible to apply all algorithms on a single extreme-scale computer. This is also evident when considering the massive (including heterogeneous and streaming) data-intensive needs with the similarly massive computed-data needs—in some cases, there will be a compromise "middle-way," but certainly not in others. Similarly, we should expect different needs and strategies for discrete versus continuous data and models.

Collaboration

The great scope of DOE-wide investments, skills, and responsibilities in national security broadly have created an environment and expertise within which to reach for the new level of "co-design" envisioned. New modes of collaboration within and between national laboratories, universities, and industry will be essential. Some of these collaborations are already being discussed in the research community but they need to be empowered and accelerated. The crosscutting opportunities across all of DOE are numerous, as this workshop revealed, and cooperation and planning between the National Nuclear Security Administration, the DOE Office of Science, and Applied Energy Programs have a high potential for maximizing impact.

Training the Next Generation of Multidisciplinary Scientists

Training the next generation of discoverers, designers, and engineers to work in true multidisciplinary co-design environments must receive every encouragement—for example, dedicated programs for training in national laboratories, but in multidisciplinary teams and supported by graduate curricula at universities. Again, there are good exploratory models at small scales in the current community but these need scaling and empowerment of a higher order. These opportunities should also be available to researchers at various stages in their careers—not only to undergraduate and graduate students. This underscores that the urgency of the challenges and the timeliness of the opportunities means no talent should be wasted.

REFERENCES

Abraham NL and MIJ Probert. 2006. "A Periodic Genetic Algorithm with Real-Space Representation for Crystal Structure and Polymorph Prediction." *Physical Review B* 73(22): Art. No. 224104. DOI: 10.1103/PhysRevB.73.224104.

Bader D, C Covey, W Gutowski, I Held, K Kunkel, R Miller, R Tokmakian, and M Zhang. 2008. *Climate Models: An Assessment of Strengths and Limitations: Synthesis and Assessment Product 3.1 Report by the U.S. Climate Change Science Program and the Subcommittee on Global Change Research.* U.S. Climate Change Science Program, Washington, D.C. Accessed March 22, 2010, at http://www.climatescience.gov/Library/sap/sap3-1/final-report/sap3-1-final-all.pdf.

Baran A, A Staszczak, J Dobaczewski, and W Nazarewicz. 2007. "Collective Inertia and Fission Barriers within the Skyrme-Hartree-Fock Theory." *International Journal of Modern Physics E (IJMPE)* 16(2):443-451. DOI: 10.1142/S0218301307005879.

Batcheller TA and GD Redden. 2004. *Colloidal Plutonium at the OU 7-13/14 Subsurface Disposal Area: Estimate of Inventory and Transport Properties*. Idaho Completion Project (ICP/EXT-04-00253), Bechtel BWXT Idaho LLC, Idaho Falls, Idaho.

Beane SR, PF Bedaque, K Orginos, and MJ Savage. 2006. "Nucleon-Nucleon Scattering from Fully-Dynamical Lattice QCD." *Physical Review Letters* 97(1):012001. DOI:10.1103/PhysRevLett.97.012001.

Beane SR, K Orginos, and MJ Savage. 2008. "Hadronic Interactions from Lattice QCD." *International Journal of Modern Physics E (IJMPE)* 17(7):1157-1218. DOI:10.1142/S0218301308010404.

Behler J, R Martoňák, D Donadio, and M Parrinello. 2008a. "Metadynamics Simulations of the High-Pressure Phases of Silicon Employing a High-Dimensional Neural Network Potential." *Physical Review Letters* 100(18):185501. DOI:10.1103/PhysRevLett.100.185501.

Behler J, R Martoňák, D Donadio, and M Parrinello. 2008b. "Pressure-Induced Phase Transitions in Silicon Studied by Neural Network-Based Metadynamics Simulations." *Physica Status Solidi B* 245(12):2618-2629. DOI: 10.1002/pssb.200844219.

Berger JF, M Girod, and D Gogny. 1984. "Microscopic Analysis of Collective Dynamics in Low Energy Fission." *Nuclear Physics A* 428:23-36. DOI: 10.1016/0375-9474(84)90240-9.

Berghout HL, SF Son, LG Hill, and BW Asay. 2006. "Flame Spread Through Cracks of PBX 9501 (A Composite Octahydro-1,3,5,7-tetranitro-1,3,5,7-tetrazocine-Based Explosive)." *Journal of Applied Physics* 99(11):114901. DOI: 10.1063/1.2196219.

Bertsch GF, DJ Dean, and W Nazarewicz. 2007. "Universal Nuclear Energy Density Functional: Computing Atomic Nuclei." *SciDAC Review* 6:42-51.

Brecht SH, DW Hewett and DJ Larson. 2009. "A Magnetized, Spherical Plasma Expansion in an Inhomogeneous Plasma: Transition from Super- to Sub-Alfvénic," *Geophysical Research Letters* 36(15): Article No.: L15105. DOI:10.1029/2009GL038393.

Bromberg JL. 1982. *Fusion: Science, Politics, and the Invention of a New Energy Source.* The MIT Press, Cambridge, Massachusetts.

Brooks HE and CA Doswell. 1996. "A Comparison of Measures Oriented and Distributions-Oriented Approaches to Forecast Verification." *Weather and Forecasting* 11:288-303.

Brown WM, AP Thompson, and PA Schultz. 2010. "Efficient Hybrid Evolutionary Optimization of Interatomic Potential Models." *Journal of Chemical Physics* 132(2): 024108. DOI: 10.1063/1.3294562.

Camden JP, R Dawes, and DL Thompson. 2009. "Application of Interpolating Moving Least Squares (IMLS) Fitting to Hypervelocity Collision Dynamics: $O(^3P)$ + HCl. " *Journal of Physical Chemistry A* 113(16):4626-4630. DOI: 10.1021/jp8113144.

Cavazzoni C, GL Chiarotti, S Scandolo, E Tosatti, M Bernasconi, and M Parrinello. 1999. "Superionic and Metallic States of Water and Ammonia at Giant Planet Conditions." *Science* 283(5398):44-46. DOI: 10.1126/science.283.5398.44.

Cawkwell MJ, TD Sewell, Zheng LQ, and DLThompson. 2008. "Shock-Induced Shear Bands in an Energetic Molecular Crystal: Application of Shock-Front Absorbing Boundary Conditions to Molecular Dynamics Simulations." *Physical Review B* 78(1):014107. DOI: 10.1103/PhysRevB.78.014107.

Chadola V, A Bangerjee, and V Kumar. 2009. "Anomaly Detection: A Survey." *ACM Computing Surveys* 41(3):1-72. DOI: 10.1145/1541880.1541882.

Chadwick MB, P Talou, and T Kawano. 2005. "Reducing Uncertainties in Nuclear Data." *Los Alamos Science* 29:26-41. Accessed March 22, 2010, at http://library.lanl.gov/cgi-bin/getfile?29-03.pdf.

Chadwick MB, T Kawano, P Talou, E Bauge, S Hilaire, P Dossantos-Uzarralde, PE Garrett, JA Becker, and RO Nelson. 2007. "Yttrium ENDF/B-VII Data from Theory and LANSCE/GEANIE Measurements and Covariances Estimated Using Bayesian and Monte-Carlo Methods." *Nuclear Data Sheets* 108(12):2742-2755. DOI: 10.1016/j.nds.2007.11.004.

Chapman DR. 1979. "Computational Aerodynamics Development and Outlook." *American Institute of Aeronautics and Astronautics Journal* 17(12):1293-1313.

Chen YC, K Nomura, RK Kalia, A Nakano, and P Vashishta. 2009. "Void Formation and Breakup in Shearing Silica Glass." *Physical Review Letters* 103(3):035501-1-035501-4. DOI:10.1103/PhysRevLett.103.035501.

Chokshi AH, A Rosen, J Karch, and H Gleiter. 1989. "On The Validity of the Hall-Petch Relationship in Nanocrystalline Materials." *Scripta Materialia* 23(10):1679-1683. DOI:10.1016/0036-9748(89)90342-6.

Council on Competitiveness. 2008a. *Reveal: Council on Competitiveness and USC-ISI Broad Study of Desktop Technical Computing End Users and HPC.* Council on Competitiveness, Washington, D.C.

Council on Competitiveness. 2008b. *Reflect: Council on Competitiveness and USC-ISI IN-Depth Study of Technical Computing End Users and HPC.* Council on Competitiveness, Washington, D.C.

Ciccotti M. 2009. "Stress-Corrosion Mechanisms in Silicate Glasses." *Journal of Physics D* 42(21):214006. DOI:10.1088/0022-3727/42/21/214006.

Dawes R, AF Wagner, and DL Thompson. 2009. "Ab Initio Wavenumber Accurate Spectroscopy: $^{1}CH_2$ and HCN Vibrational Levels on Automatically Generated IMLS Potential Energy Surfaces." *Journal of Physical Chemistry A* 113(16):4709-4721. DOI: 10.1021/jp900409r.

Deumens E, A Diz, R Londo, and Y Öhrn. 1994. "Time-Dependent Theoretical Treatments of the Dynamics of Electrons and Nuclei in Molecular Systems." *Review of Modern Physics* 66(3):917-983. DOI: 10.1103/RevModPhys.66.917.

DOE. 2006. *Basic Research Needs for Advanced Nuclear Energy Systems*. U.S. Department of Energy, Washington, D.C.

DOE. 2008. *Basic Research Needs for Materials Under Extreme Environments*. U.S. Department of Energy, Washington, D.C.

DOE. 2009. *Next Generation Photon Sources for Grand Challenges in Science and Energy*. U.S. Department of Energy and Basic Energy Sciences Advisory Committee, Washington, D.C. Accessed April 1, 2010, at http://www.science.doe.gov/bes/reports/abstracts.html#NGPS.

Donev A, VV Butatov, T Oppelstrup, GH Gilmer, B Sadigh, and MH Kalos. 2010. "First-Passage Kinetic Monte Carlo Algorithm for Complex Diffusion-Reaction Systems." *Journal of Computational Physics* 229(9):3214-3236. DOI:10.1016/j.jcp.2009.12.038.

Garrett, M. 2006. *Before the United States Senate Committee on Commerce, Science and Transportation Subcommittee on Technology, Innovation, and Competitiveness*. Testimony of Michael Garrett Director, Airplane Performance Boeing Commercial Airplanes, July 19, 2006.

Gaffney FJ, Jr. 2005. "EMP: America's Achilles' Heel." *The National Speech Digest of Hillsdale College Imprimis* 34(6):1-5.

Germann TC and K Kadau. 2008. "Trillion Atom Molecular Dynamics Simulation Becomes a Reality." *International Journal of Modern Physics C* 19(9): 1315-1319. DOI: 10.1142/S0129183108012911.

Gindensperger E, C Meier, and JA Beswick. 2000. "Mixing Quantum and Classical Dynamics Using Bohmian Trajectories." *Journal of Chemical Physics* 113(21):9369-9372. DOI: 10.1063/1.1328759.

Glasstone S and PJ Dolan. 1977. *The Effects of Nuclear Weapons*. 3rd ed. U.S. Government Printing Office, Washington, D.C.

Goldman N, LE Fried, IFW Kuo, and CJ Mundy. 2005. "Bonding in the Superionic Phase of Water." *Physical Review Letters* 94(21):217801. DOI: 10.1103/PhysRevLett.94.217801.

Goutte H, JF Berger, P Casoli, and D Gogny. 2005. "Microscopic Approach of Fission Dynamics Applied to Fragment Kinetic Energy and Mass Distributions in ^{238}U." *Physical Review C* 71(2): Art. No. 024316. DOI: 10.1103/PhysRevC.71.024316.

Hagen G, DJ Dean, M Hjorth-Jensen, and T Papenbrock. 2007a. "Complex Coupled-Cluster Approach to an *Ab-Initio* Description of Open Quantum Systems." *Physics Letters B* 656(4-5):169-173. DOI:10.1016/j.physletb.2007.07.072.

Hagen G, DJ Dean, M Hjorth-Jensen, T Papenbrock, and A Schwenk. 2007b. "Benchmark Calculations for ^{3}H, ^{4}He, ^{16}O, and ^{40}Ca with *Ab Initio* Coupled-Cluster Theory." *Physical Review C* 76(4):044305. DOI: 10.1103/PhysRevC.76.044305.

Hart GLW, V Blum, MJ. Walorski, and A Zunger. 2005. "Evolutionary Approach for Determining First-Principles Hamiltonians." *Nature Materials* 4(5):391-394. DOI:10.1038/nmat1374.

Hattar K, DM Follstaedt, JA Knapp, and IM Robertson. 2008. "Defect Structures Created During Abnormal Grain Growth in Pulsed-Laser Deposited Nickel." *Acta Materialia* 56(4):794-801. DOI:10.1016/j.actamat.2007.10.027.

Horsfield AP, DR Bowler, AJ Fisher, TN Todorov, and CG Sanchez. 2005. "Correlated Electron-Ion Dynamics: The Excitation of Atomic Motion by Energetic Electrons." *Journal of Physics: Condensed Matter* 17(30): 4793-4812.

Hosemann P, M Hawley, D Koury, JG Swadener, J Welch, AL Johnson, G Mori, and N Li. "Characterization of Oxide Layers Grown on D9 Austenitic Stainless Steel in Lead Bismuth Eutectic." 2008. *Journal of Nuclear Materials* 375(3):323-330.

Jeffreys H. 1967. Statistical Methods in Seismology." In: *International Dictionary of Geophysics*, K Runcorn, Editor, Pergamon, London, pp. 1398-1401.

Jordan GC IV, RT Fisher, DM Townsley, AC Calder, C Graziani, S Asida, DQ Lamb, and JW Truran. 2008. "Three-Dimensional Simulations of the Deflagration Phase of the Gravitationally Confined Detonation Model of Type Ia Supernovae." *Astrophysical Journal* 681(2):1448-1457.

Kawano T, P Talou, MB Chadwick, and T Watanabe. 2010. "Monte Carlo Simulation for Particle and γ-Ray Emissions in Statistical Hauser-Feshbach Model." *Journal of Nuclear Science and Technology* 47(5): 462-469.

Kennedy BT. November 24, 2008. "What a Single Nuclear Warhead Could Do." *The Wall Street Journal*, p. A19.

Kim WW and S Syed. 2004. "Large-Eddy Simulation Needs for Gas Turbine Combustor Design." In *Proceedings of the 42nd AIAA Aerospace Sciences Meeting and Exhibit*, pp. 3162-3175, American Institute of Aeronautics and Astronautics Inc., Reston, Virginia.

Knudson MD, MP Desjarlais, and DH Dolan. 2008. "Shock-Wave Exploration of the High-Pressure Phases of Carbon." *Science* 322(5909):1822-1825. DOI: 10.1126/science.1165278.

Koch CC and J Narayan. 2000. "The Inverse Hall-Petch Effect – Fact or Artifact?" Symposium B: Structure and Mechanical Properties of Nanophase Materials-Theory and Computer Simulations vs. Experiment. *MRS Symposia Proceedings No. 634*, Paper No. B5.1, Materials Research Society, Pittsburgh, Philadelphia.

Korniss G, MA Novotny, H Guclu, Z Totoczkai, and PA Rikvold. 2003. "Suppressing Roughness in Virtual Times in Parallel Discrete-Event Simulations." *Science* 299(5607):677-679. DOI:10.1126/science.1079382.

Krasheninnikov AV, Y Miyamoto, and D Tománek. 2007. "Role of Electronic Excitations in Ion Collisions with Carbon Nanostructures." *Physical Review Letters* 99(1): Art. No. 016104. DOI:10.1103/PhysRevLett.99.016104.

Kumar V, A Lazarevic, and J Srivastava, Editors. 2005. *Managing Cyber Threats: Issues, Approaches and Challenges*. Springer, New York.

Kushima A, X Lin, J Li, X Qian, J Eapen, JC Mauro, P Diep, and S Yip. 2009. "Computing the Viscosity of Supercooled Liquids." *Journal of Chemical Physics* 130(22): 224504-224504-12.

Kusnezov D. 2004. *Advanced Simulation and Computing—The Next Ten Years*. NA-ASC-100R-04-Vol.1, Rev. 0 (SAND 2004-3740P), Sandia National Laboratories, Albuquerque, New Mexico.

Foster, JS Jr, E Gjelde, WR Graham, RJ Hermann, HM Kluepfel, RL Lawson, GK Soper LL Wood, Jr. and JB Woodard. 2004. *Report of the Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack, Volume 1: Executive Report*. Defense Technical Information Center, Ft. Belvoir, Virginia. Accessed April 27, 2010 at http://www.empcommission.org/docs/empc_exec_rpt.pdf.

Laio A and M Parrinello. 2002. "Escaping Free-Energy Minima." *Proceedings of the National Academy of Sciences* 99:12562.

Laio A and FL Gervasio. 2008. "Metadynamics: A Method to Simulate Rare Events and Reconstruct the Free Energy in Biophysics, Chemistry and Material Science." *Reports on Progress in Physics* 71:126601.

Lau TT, A Kushima, and S Yip. 2010. "Atomistic Simulation of Creep in a Nanocrystal." *Physical Review Letter* 104(17):175501.

Lee B-J and MI Baskes. 2000. "Second Nearest-Neighbor Modified Embedded-Tom-Method Potential." *Physical Review B* 62(13):8564-8567. DOI:10.1103/PhysRevB.62.8564.

Leonard RA, GF Vandegrift, DG Kalina, DF Fischer, RW Bane, L Burris, EP Horwitz, R Chiarizia, and H Diamond. 1985. "Extraction and Recovery of Plutonium and Americium from Nitric Acid Waste Solutions by the TRUEX Process - Continuing Development Studies." ANL-85-45, Argonne National Laboratory, Argonne, Illinois. DOI: 10.2172/5160292.

Lindzen RS and YS Choi. 2009. "On the Determination of Climate Feedbacks from ERBE Data." *Geophysical Research Letters* 36(16): Art. No. L16705. DOI:10.1029/2009GL039628.

Lootens D, P Hebraud, E Lecolier, and H Van Damme. 2004. "Gelation, Shear-Thinning, Shear-Thickening in Cement Slurries." *Oil & Gas Science and Technology – Revue de I'IFP* 59(1):31-40.

Loreti CP, SA Foster, and JE Obbagy. 2001. *An Overview of Greenhouse Gas Emissions Verification Issues.* Arthur D. Little, Inc., Cambridge, Massachusetts.

Luo SN, Q An, TC Germann TC, and LB Han. 2009. "Shock-Induced Spall in Solid and Liquid Cu at Extreme Strain Rates." *Journal of Applied Physics* 106(1):013502. DOI: 10.1063/1.3158062.

Luu T and A Walker-Loud. 2009. "Nuclear Forces and High-Performance Computing: The Perfect Match." *Journal of Physics: Conference Series* 180:012081. DOI: 10.1088/1742-6596/180/1/012081.

Martoňák R, A Laio, and M Parrinello. 2003. "Predicting Crystal Structures: The Parrinello-Rahman Method Revisited." *Physical Review Letters* 90:75503.

Mattsson AE, R Armiento, J Paier, G Kresse, JM Wills, and TR Mattsson. 2008. "The AM05 Density Functional Applied to Solids." *Journal of Chemical Physics* 128(8):084714.

Mattsson TR and MP Desjarlais. 2006. "Phase Diagram and Electrical Conductivity of High Energy-Density Water from Density Functional Theory." *Physical Review Letters* 97(1):017801. DOI: 10.1103/PhysRevLett.97.017801.

McAvaney BJ, C Covey, S Joussaume, V Kattsov, A Kitoh, W Ogana, AJ Pitman, AJ Weaver, RA Wood, ZC Zhao, J AchutaRao, A Arking, A Barnston, R Betts, C Bitz, G Boer, P Broccoli, F Bryan, M Claussen, R Colman, P Delecluse, A Del Genio, K Dixon, P Duffy, L DüMenil, M England, T Fichefet, G Flato, JC Fyfe, N Gedney, P Gent, C Genthon, J Gregory, E Guilyardi, S Harrison, N Hasegawa, G Holland, Y Jia, PD Jones, K Kageyama, D Keith, K Kodera, J Kutzbach, S Lambert, S Legutke, G Madec, S Maeda, ME Mann, G Meehl, I Mokhov, T Motoi, T Phillips, J Polcher, GL Potter, V Pope, C Prentice, G Roff, F Semazzi, P Sellers, DJ Stensrud, T Stockdale, R Stouffer, KE Taylor, K Trenberth, R Tol, J Walsh, M Wild, D Williamson, SP Xie, XH Zhang, F Zwiers. 2001. "Model Evaluation." Chapter 8 in *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change,* eds. JT Houghton, Y Ding, DJ Griggs, M Noguer, PJ van der Linden, X Dai, K Maskell and CA Johnson. Cambridge University Press, Cambridge, United Kingdom.

McIntyre DH and BB Tussing. 2008. "Threats at our Threshold: Homeland Defense and Homeland Security in the New Century." Compilation of Conference Proceedings, U.S. Army War College, August 18, 2008.

McGrane SD, A Grieco, KJ Ramos, DE Hooks, and DS Moore. 2009. "Femtosecond Micromachining of Internal Voids in High Explosive Crystals for Studies of Hot Spot Initiation." *Journal of Applied Physics* 105(7):073505.

Meehl GA, TF Stocker, WD Collins, P Friedlingstein, AT Gaye, JM Gregory, A Kitoh, R Knutti, JM Murphy, A Noda, SCB Raper, IG Watterson, AJ Weaver ZC Zhao, RB Alley, J Annan, J Arblaster, C Bitz, P Brockmann, V Brovkin, L Buja, P Cadule, G Clarke, M Collier, M Collins, E Driesschaert, NA Diansky, M Dix, K Dixon, JL Dufresne, M Dyurgerov, M Eby, NR Edwards, S Emori, P Forster, R Furrer, P Gleckler, J Hansen, G Harris, GC Hegerl, M Holland, A Hu, P Huybrechts, C Jones, F Joos, JH Jungclaus, J Kettleborough, M Kimoto, T Knutson, M Krynytzky, D Lawrence, A Le Brocq, MF Loutre, J Lowe, HD Matthews, M Meinshausen, SA Müller, S Nawrath, J Oerlemans, M Oppenheimer, J Orr, J Overpeck, T Palmer, A Payne, GK Plattner, J Räisänen, A Rinke, E Roeckner, GL Russell, D Salas y Melia, B Santer, G Schmidt, A Schmittner, B Schneider, A Shepherd, A Sokolov, D Stainforth, PA Stott, RJ Stouffer, KE Taylor, C Tebaldi, H Teng, L Terray, R van de Wal, D Vaughan, EM Volodin, B Wang, TML Wigley, M Wild, J Yoshimura, R Yu, and S Yukimoto. 2007. "Global Climate Projections." Chapter 10 in *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, S Solomon, D Qin, M Manning, M Marquis, K Avery, MMB Tignor, HL Miller, and Z Chen, Editors. Cambridge University Press, Cambridge, United Kingdom.

Menikoff R and MS Shaw. 2008. "Review of the Forest Fire Model for High Explosives." *Combustion Theory and Modeling* 12(3):569-604. DOI: 10.1080/13647830801942402.

Moin P and J Kim. 1997. "Tackling Turbulence with Supercomputers." *Scientific American* 276(1):62-68.

Nakano A, RK Kalia, KI Nomura, A Sharma, P Vashishta, F Shimojo, ACT Van Duin, WA Goddard, R Biswas, D Srivastava, and LH Yang. 2008. "De Novo Ultrascale Atomistic Simulations on High-End Parallel Supercomputers." *International Journal of High Performance Computing Applications* 22(1):113-128. DOI: 10.1177/1094342007085015.

National Research Council. 1979. *Carbon Dioxide and Climate: A Scientific Assessment*. Ad Hoc Study Group on Carbon Dioxide and Climate. The National Academy of Sciences, Washington D.C.

National Research Council. 2003. *Frontiers in High Energy Density Physics: The X-Games of Contemporary Science*. The National Academies Press, Washington D.C.

Navrátil P, JP Vary, and BR Barrett. 2000. "Large-Basis *Ab Initio* No-Core Shell Model and its Application to ^{12}C." *Physical Review C* 62(5):054311. DOI: 10.1103/PhysRevC.62.054311.

Navrátil P, VG Gueorguiev, JP Vary, WE Ormand, and A Nogga. 2007. "Structure of A=10–13 Nuclei with Two- Plus Three-Nucleon Interactions from Chiral Effective Field Theory." *Physical Review Letters* 99(4): Art. No. 042501. DOI: 10.1103/PhysRevLett.99.042501.

Negele JW and H Orland. 1987. *Quantum Many-Particle Systems (Frontiers in Physics)*. Perseus Books Group, New York.

Netzloff HM, MA Collins, and MS Gordon. 2006. "Growing Multiconfigurational Potential Energy Surfaces with Applications to $X+H_2$ (X=C,N,O) Reactions." *Journal of Chemical Physics* 124(15):154104. DOI: 10.1063/1.2185641.

Nollett KM, SC Pieper, RB Wiringa, J Carlson, and GM Hale. 2007. "Quantum Monte Carlo Calculations of Neutron-α Scattering." *Physical Review Letters* 99(2): Art. No. 022502. DOI: 10.1103/PhysRevLett.99.022502.

Nomura K, H Dursun, R Seymour, W Wang, RK Kalia, A Nakano, P Vashishta, F Shimojo, and LH Yang. 2009. "A Metascalable Computing Framework for Large Spatiotemporal-Scale Atomistic Simulations." *Proceedings of the International Parallel and Distributed Processing Symposium*. Institute of Electrical and Electronics Engineers, New York.

Northrop JA. 1996. *Handbook of Nuclear Weapon Effects: Calculational Tools Abstracted from EM-1*. 1st ed. Defense Special Weapons Agency, Alexandria, Virginia.

Novikov AP, SN Kalmykov, S Utsunomiya, RC Ewing, F Horreard, A Merkulov, SB Clark, VV Tkachev, and BF Myasoedov. 2006. "Colloid Transport of Plutonium in the Far-Field of the Mayak Production Association, Russia." *Science* 314(5799):638-641. DOI: 10.1126/science.1131307.

Nowick AS and BS Berry. 1972. *Anelastic Relaxation in Crystalline Solids*. Academic Press, New York.

Oganov AR and CW Glass. 2006. "Crystal Structure Prediction Using *Ab Initio* Evolutionary Techniques: Principles and Applications." *Journal of Chemical Physics* 124(24):244704.

Olmsted D, SM Foiles, and EA Holm. 2009a. "Survey of Computed Grain Boundary Properties in Face-Centered Cubic Metals: I. Grain Boundary Energy." *Acta Materialia* 57:3694-3703.

Olmsted D, EA Holm, and SM Foiles. 2009b. "Survey of Computed Grain Boundary Properties in Face-Centered Cubic Metals: II. Grain Boundary Mobility." *Acta Materialia* 57:3704-3713.

Oppelstrup T, VV Bulatov, GH Gilmer, MH Kalos, and B Sadigh. 2006. "First-Passage Monte Carlo Algorithm: Diffusion Without All the Hops." *Physical Review Letters* 97:230602.

Oppelstrup T, VV Bulatov, A Donev, MH Kalos, GH Gilmer, and B Sadigh. 2009. "First-Passage Kinetic Monte Carlo Method." *Physical Review E* 80(6): Art. No. 066701.

Pathania R. 2005. *Status Review of Initiation of Environmentally Assisted Cracking and Short Crack Growth*. Electric Power Research Institute (EPRI) Report 1011788, Palo Alto, California.

Pei JC, W Nazarewicz, JA Sheikh, and AK Kerman. 2009. "Fission Barriers of Compound Superheavy Nuclei." *Physical Review Letters* 102(19):192501. DOI: 10.1103/PhysRevLett.102.192501.

Pieper SC and RB Wiringa. 2001. "Quantum Monte Carlo Calculations of Light Nuclei." *Annual Review of Nuclear and Particle Science* 51:53-90. DOI: 10.1146/annurev.nucl.51.101701.132506.

Peterson, VL. 1989. "Computational Challenges in Aerospace." *Future Generation Computer Systems*, 5(2-3): 243-258.

Quaglioni S and P Navrátil. 2008. "*Ab Initio* Many-Body Calculations of n-^{3}H, n-^{4}*He*, p-^{3}, ^{4}*He*, and n-^{10}*Be* Scattering." *Physical Review Letters* 101(9):092501. DOI: 10.1103/PhysRevLett.101.092501.

Reynolds WC, JJ Alonso and M Fatica, 2003. "Aircraft Gas Turbine Engine Simulations." In *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference,* Paper No.: 2003-3698. June 23-26, 2003, Orlando, Florida. American Institute of Aeronautics and Astronautics Inc., Reston, Virginia.

Reichman DR, P-N Roy, S Jang, and GA Voth. 2000. "A Feynman Path Centroid Dynamics Approach for the Computation of Time Correlation Functions Involving Nonlinear Operators." *Journal of Chemical Physics* 113(3):919-929. DOI: 10.1063/1.481872.

Roberto J and TD de La Rubia. 2006. *Report of the Basic Energy Sciences Workshop on Basic Research Needs for Advanced Nuclear Energy Systems*. U.S. Department of Energy, Office of Basic Energy Sciences, Washington D.C.

Santer DB, TML Wigley, ME Schlesinger, and JFB Mitchell. 1990. *Developing Climate Scenarios from Equilibrium GCM Results*. Report No. 47, Max-Planck-Institut für Meteorologie, Hamburg, Germany.

Schiøtz J, FD DiTolla, and KW Jacobsen. 1998. "Softening of Nanocrystalline Metals at Very Small Grain Sizes." *Nature* 391:561-563.

Schiøtz J and KW Jacobsen. 2003. "A Maximum in the Strength of Nanocrystalline Copper." *Science* 301(5638):1357-1359. DOI: 10.1126/science.1086636.

Sheikh JA, W Nazarewicz, and JC Pei. 2009. "Systematic Study of Fission Barriers of Excited Superheavy Nuclei." *Physical Review C* 80(1): Art. No. 011302(R). DOI:10.1103/PhysRevC.80.011302.

Schneider M. 2007. *The Emerging EMP Threat to the United States.* National Institute Press, Fairfax, Virginia.

Skalski J. 2007. "Nuclear Fission with Mean-Field Instantons." *Physical Review C* 77(6): Art. No. 064610. DOI: 10.1103/PhysRevC.77.064610.

Slepoy A, MD Peters, and AP Thompson. 2007. "Searching for Globally Optimal Functional Forms for Interatomic Potentials Using Genetic Programming with Parallel Tempering." *Journal of Computational Chemistry* 28(15):2465-2471.

Smilowitz L, BF Henson, JJ Romero, BW Asay, CL Schwartz, A Saunders, FE Merrill, CL Morris, K Kwiatkowski, G Hogan, P Nedrow, MM Murray, TN Thompson, W McNeil, P Rightley, and M Marr-Lyon. 2008. "Direct Observation of the Phenomenology of a Solid Thermal Explosion Using Time-Resolved Proton Radiography." *Physical Review Letters* 100(22): Art. No. 228301. DOI:10.1103/PhysRevLett.100.228301.

Soderholm L, PM Almond, S Skanthakumar, RE Wilson, and PC Burns. 2008. "The Structure of the Plutonium Oxide Nanocluster $[Pu_{38}O_{56}Cl_{54}(H_2O)_8]^{14-}$." *Angewandte Chemie International Edition* 47(2):298-302. DOI: 10.1002/anie.200704420.

Srinivasan SG and ML Baskes. 2007. "Atomistic Simulations of the Plasticity Behavior of Shock-Induced Polycrystalline Nickel." *Metallurgical and Materials Transactions* 38A:2716-2720.

Staszczak A, A Baran, J Dobaczewski, and W Nazarewicz. 2009. "Microscopic Description of Complex Nuclear Decay: Multimodal Fission." *Physical Review C* 80(1):014309. DOI: 10.1103/PhysRevC.80.014309.

Stella L, M Meister, AJ Fisher, and AP Horsfield. 2007. "Robust Nonadiabatic Molecular Dynamics for Metals and Insulators." *Journal of Chemical Physics* 127: Art. No. 214104. DOI:10.1063/1.2801537.

Stuart SJ, AB Tutein, and JA Harrison. 2000. "A Reactive Potential for Hydrocarbons with Intermolecular Interactions." *Journal of Chemical Physics* 112(14):6472. DOI: 10.1063/1.481208.

Sugita Y and O Yuko. 1999. "Replica-Exchange Molecular Dynamics Method for Protein Folding." *Chemical Physics Letters* 314(1-2):141-151.

Takemiya H, Y Tanaka, S Sekiguchi, S Ogata, RK Kalia, A Nakano, and P Vashishta. 2006. "Sustainable Adaptive Grid Supercomputing: Multiscale Simulation of Semiconductor Processing Across the Pacific." *Proceedings of Supercomputing 2006*, Institute of Electrical and Electronics Engineers/American Chemical Society, New York.

Thomson A. February 20-26, 1995a. "Civil Satellite Vulnerability." *Space News;* p.15.

Thomson A. 1995b. "Satellite Vulnerability: A View from the USAF." Message from Allen Thomson to Internet newsgroups sci.space.policy,alt.war,alt.politics.org.cia posted November 22, 1995. Accessed April 22, 2010 at http://www.fas.org/spp/military/program/asat/at_951122.htm.

(Top 2010) TOP500.Org. 2009. *Top 500 Supercomputer Sites*. Accessed June 9, 2010, at http:/www.top500.org/. Last updated April 28, 2010.

Tully JC. 1990. "Molecular Dynamics with Electronic Transitions." *Journal of Chemical Physics* 93(2):1061. DOI:10.1063/1.459170.

van Duin ACT, S Dasgupta, F Lorant, and WA Goddard III. 2001. "ReaxFF: A Reactive Force Field for Hydrocarbons." *Journal of Physical Chemistry A* 105(41):9396-9409. DOI: 10.1021/jp004368u.

Van Swygenhoven H and A Caro. 1997. "Molecular Dynamics Computer Simulation of Nanophase Ni: Structure and Mechanical Properties." *Nanostructured Materials* 9(1-8):669-672. DOI:10.1016/S0965-9773(97)00147-5.

Vashishta P, RK Kalia, A Nakano, E Kaxiras, A Grama, G Lu, S Eidenbenz, AF Voter, and RQ Hood. 2008. JA Moriarty and LH Yang, Eds. "Hierarchical Petascale Simulation Framework for Stress Corrosion Cracking." *Journal of Physics: Conference Series* 125: Art. No. 012060. DOI: 10.1088/1742-6596/125/1/012060.

Veynante D and L Vervisch. 2002. "Turbulent Combustion Modeling." *Progress in Energy and Combustion Science* 28(3):193-266.

Vo NQ, RS Averback, P Bellon, and A Caro. 2008. "Limits of Hardness at the Nanoscale: Molecular Dynamics Simulations." *Physical Review B* (Rapid Communications) 78: Art. No. 241402.

Voter AF. 1998. "Parallel Replica Method for Dynamics of Infrequent Events." *Physical Review B* 57(22):13985-13988. DOI: 10.1103/PhysRevB.57.R13985.

Weiss GM. 2004. "Mining with Rarity: A Unifying Framework." *Special Interest Group on Knowledge Discovery and Data Mining: Explorations* 6(1):7-19. Accessed March 22, 2010, at http://www.er.doe.gov/bes/reports/files/MUEE_rpt.pdf.

Wu CJ, LE Fried, LH Yang, N Goldman, and S Bastea. 2009. "Catalytic Behaviour of Dense Hot Water." *Nature Chemistry* 1:57-62. DOI: 10.1038/nchem.130.

Yildiz B, A Nikiforova, and S Yip. 2009. "Metallic Interfaces in Harsh Chemo-Mechanical Environments." *Nuclear Engineering and Technology* 41(1):21.

Yip S. 2003. "Synergistic Science." *Nature Materials* 2(1):3-5. DOI:10.1038/nmat778.

Yip S. 2009. "Multiscale Materials." In: *Multiscale Methods*, J Fish, Ed., Oxford University Press, Chapter 14.

Younes W and D Gogny. 2009a. "Microscopic Calculation of ^{240}Pu Scission with a Finite-Range Effective Force." *Physical Review C* 80(5): Art. No. 054313. DOI: 10.1103/PhysRevC.80.054313.

Younes W and D Gogny. 2009b. "The Microscopic Theory of Fission." In: *Proceedings of the 4th International Workshop on Nuclear Fission and Fission-Product Spectroscopy*, Vol. 1175. A Chatillon, H Faust, G Fioni, D Goutte, and H Goutte, Eds. 4th International Workshop on Nuclear Fission and Fission Product Spectroscopy, Cadarache, France.

Zhang LZ, SV Zybin, ACT van Duin, S Dasgupta, WA Goddard III, and EM Kober. 2009. "Carbon Cluster Formation during Thermal Decomposition of Octahydro-1,3,5,7-tetranitro-1,3,5,7-tetrazocine and 1,3,5-Triamino-2,4,6-trinitrobenzene High Explosives from ReaxFF Reactive Molecular Dynamics Simulation." *Journal of Physical Chemistry A* 113(40):10619-10640. DOI: 10.1021/jp901353a.

Zhu T, J Li, and S Yip. 2005. "Stress-Dependent Molecular Pathways of Silica-Water Reaction." *Journal of Mechanics and Physics of Solids* 53:1597-1623.

APPENDICES

APPENDIX 1: WORKSHOP AGENDA

APPENDIX 2: WORKSHOP PARTICIPANTS

APPENDIX 3: ACRONYMS

APPENDIX 1: WORKSHOP AGENDA

Monday, October 5, 2009

Time	Session	Lead	Room
7:00 - 8:30 p.m.	Pre-Workshop Dinner (Organizers, Panel Leads, & Speakers)		Frederick
8:30 p.m.	Adjourn		

Tuesday, October 6, 2009

Time	Session	Lead	Room
7:30-8:30 a.m.	Registration/Morning Refreshments		Foyer
	Welcome and Plenary Talks		
8:30- 8:45 a.m.	Welcome, Logistics	Alan Bishop, LANL & Paul Messina, ANL	Montgomery
8:45 - 9:15 a.m.	The View from SC- ASCR	Michael Strayer, Associate Director, Office of Science	
9:15 - 9:45 a.m.	The View from NNSA-ASC	David Crandall, Chief Scientist, NNSA	
9:45 - 10:15 a.m.	The Role of Computing in the Defense Program's "Predictive Capability Framework"	Robert Webster, Deputy Science Advisor to Defense Programs (NNSA/LANL)	
10:15 a.m.	**General Discussion**		Foyer
10:45 - 11:15 a.m.	Extreme-Scale Computing: Transformational Opportunities In Science	Jim Roberto, Director of Strategic Capabilities, ORNL	Montgomery
11:15 - 11:45 a.m.	The Imperative of High-Performance Computing: Pushing Back the Frontiers of Science and Technology	William F. Brinkman, Director, Office of Science	
11:45 - 12:00 p.m.	Stage Setting: Purpose of Panel Sessions, Review of Agenda, Plans for Report	Alan Bishop & Paul Messina	
12:00 - 1:15 p.m.	Working Lunch: Extreme Computing, The Department of Energy and the Obama Administration: A Strategic Perspective	Victor Reis, Senior Advisor, Office of the Secretary, DOE	

Time	Session	Lead	Room
1:15-3:30 p.m.	**Breakout Sessions**		
	Multiphysics Simulation Problems	Robert Rosner & Mike Zika	Rockville
	Nuclear Physics	Joseph Carlson & Edward Hartouni	Potomac
	Materials Science	John Sarrao & Sidney Yip	Darnestown
	Chemistry	Robert Harrison & Steve Plimpton	Gaithersburg
	The Science of Nonproliferation	Alex Pothen & Mark Rintoul	Frederick
	Uncertainty Quantification and Error Analysis	David Higdon & Richard Klein	Bethesda
3:30 p.m.	General Discussion		Foyer
3:45-6:00 p.m.	**Continue Breakout Sessions (return to breakout rooms)**		
6:30 - 8:00 p.m.	Working Dinner: Organizational meetings		Salons A & B
8:00 p.m.	Adjourn		

Wednesday, October 7, 2009

Time	Session	Lead	Room
7:30-8:30 a.m.	Morning Refreshments & Preparation for Breakouts		Foyer
8:30 a.m.-10:00 a.m.	**Breakout Sessions**		
	Multiphysics Simulation Problems	Robert Rosner & Mike Zika	Rockville
	Nuclear Physics	Joseph Carlson & Edward Hartouni	Potomac
	Joint Session: Materials Science and Chemistry	John Sarrao & Sidney Yip; Robert Harrison & Steve Plimpton	Darnestown & Gaithersburg
	The Science of Nonproliferation	Alex Pothen & Mark Rintoul	Frederick
	Uncertainty Quantification and Error Analysis	David Higdon & Richard Klein	Bethesda
10:00 a.m.	General Discussion		Foyer
10:15 - 11:00 a.m.	**Continue Breakout Sessions** (return to breakout rooms)		
	Plenary Session		
11:00 – 11:45 a.m.	Update on Exascale Planning and Overview of Exascale Technology Issues	Rick Stevens, Argonne National Laboratory & Andrew White, Los Alamos National Laboratory	Montgomery

Time	Session	Lead	Room
11:45 - 1:00 p.m.	Working Lunch: Organizational Meetings		Foyer & Montgomery
Plenary Session			
1:00 – 2:45 p.m.	Presentations from Breakout Session Leaders		Montgomery
2:45 p.m.	General Discussion		Foyer
3:00 - 5:30 p.m.	**Continue Breakout Sessions** (return to breakout rooms)		
6:00 - 7:30 p.m.	Working Dinner: Organizational Meetings		Salons D & E
7:30 p.m.	Adjourn		

Thursday, October 8, 2009

Time	Session	Lead	Room
7:30-8:30 a.m.	Morning Refreshments		Foyer
Plenary Session			
8:30 - 10:00 a.m.	Workshop Summary from Each Panel		Montgomery
10:00 - 10:20 a.m.	General Discussion		Foyer
10:20 - 12:00 p.m.	Report Writing Session For Chairs, Panel Leads, And Writers		Montgomery
12:00 p.m.	Adjourn		

APPENDIX 2: WORKSHOP PARTICIPANTS

Last Name	First Name	Email	Affiliation
Anderson	Mark	mcanderson@lanl.gov	Los Alamos National Laboratory
Anitescu	Mihai	anitescu@mcs.anl.gov	Argonne National Laboratory
Averback	Robert	averback@illinois.edu	University of Illinois at Urbana-Champaign
Balakrishnan	Ramesh	bramesh@anl.gov	Argonne National Laboratory
Ballinger	Ron	hvymet@mit.edu	Massachusetts Institute of Technology
Bamzai	Anjuli	anjuli.bamzai@science.doe.gov	U.S. Department of Energy, Office of Science
Battaile	Corbett	ccbatta@sandia.gov	Sandia National Laboratories
Beckman	Peter	beckman@mcs.anl.gov	Argonne National Laboratory
Belak	Jim	belak1@llnl.gov	Lawrence Berkeley National Laboratory
Bergman	Aaron	aaron.bergman@science.doe.gov	U.S. Department of Energy, Office of Science
Berliner	Mark	mb@stat.osu.edu	Ohio State University
Bishop	Alan	arb@lanl.gov	Los Alamos National Laboratory
Brinkman	William	bill.brinkman@science.doe.gov	U.S. Department of Energy, Office of Science
Brown	David	dlb@llnl.gov	Lawrence Livermore National Laboratory
Budil	Kimberly	kimberly.budil@science.doe.gov	U.S. Department of Energy, Office of Science
Bulatov	Vasily	bulatov@llnl.gov	Lawrence Livermore National Laboratory
Bulgac	Aurel	bulgac@phys.washington.edu	University of Washington
Carlson	Joseph	carlson@lanl.gov	Los Alamos National Laboratory
Caro	Alfredo	caro2@llnl.gov	Lawrence Livermore National Laboratory
Challacombe	Matt	mchalla@lanl.gov	Los Alamos National Laboratory
Chatterjee	Lali	lali.chatterjee@science.doe.gov	U.S. Department of Energy, Office of Science
Chilton	Larry	larry.chilton@pnl.gov	Pacific Northwest National Laboratory
Colella	Phil	colella@hpcrdm.lbl.gov	Lawrence Berkeley National Laboratory
Collis	Scott	sscoll@sandia.gov	Sandia National Laboratories
Covey	Curtis	covey1@llnl.gov	Lawrence Livermore National Laboratory
Crandall	David	david.crandall@nnsa.doe.gov	National Nuclear Security Administration
DeJong	Bert	wibe.dejong@pnl.gov	Pacific Northwest National Laboratory
Dean	David	deandj@ornl.gov	Oak Ridge National Laboratory
Demkowicz	Michael	demkowicz@mit.edu	Massachusetts Institute of Technology
Dendy	Edward	dendy@lanl.gov	Los Alamos National Laboratory
Doebling	Scott	doebling@lanl.gov	National Nuclear Security Administration
Drake	Paul	rpdrake@umich.edu	University of Michigan
Foiles	Stephen	foiles@sandia.gov	Sandia National Laboratories
Frazier	Njema	Njema.Frazier@nnsa.doe.gov	National Nuclear Security Administration
Fried	Larry	lfried@llnl.gov	Lawrence Livermore National Laboratory
Galli	Giulia	gagalli@ucdavis.edu	University of California, Davis
Germann	Timothy	tcg@lanl.gov	Los Alamos National Laboratory
Ghattas	Omar	omar@ices.utexas.edu	University of Texas
Goldstein	William	goldstein3@llnl.gov	Lawrence Livermore National Laboratory
Graves	Sara	sgraves@itsc.uah.edu	University of Alabama
Graziani	Carlo	carlo@oddjob.uchicago.edu	University of Chicago
Gupta	Yogi	ymgupta@wsu.edu	Washington State University
Habib	Salman	habib@lanl.gov	Los Alamos National Laboratory
Harrison	Robert	harrisonrj@ornl.gov	Oak Ridge National Laboratory
Hartouni	Edward	hartouni2@llnl.gov	Lawrence Livermore National Laboratory

Last Name	First Name	Email	Affiliation
Hatton	Angela	angela.hatton@pnl.gov	Pacific Northwest National Laboratory
Haxton	Wick	haxton@phys.washington.edu	University of California, Berkeley / University of Washington
Hayes-Sterbenz	Anna	anna_hayes@lanl.gov	Los Alamos National Laboratory
Helland	Barbara	barbara.helland@science.doe.gov	U.S. Department of Energy, Office of Science
Hendrickson	Bruce	bahendr@sandia.gov	Sandia National Laboratories
Henning	Paul	phenning@lanl.gov	Los Alamos National Laboratory
Hereld	Mark	Hereld@mcs.anl.gov	Argonne National Laboratory
Higdon	David	dhigdon@lanl.gov	Los Alamos National Laboratory
Hoang	Thuc	Thuc.Hoang@nnsa.doe.gov	National Nuclear Security Administration
Hovland	Paul	hovland@mcs.anl.gov	Argonne National Laboratory
Jameson	Leland	ljameson@nsf.gov	National Science Foundation
Johnson	Gary	GMJ@ComputionalScienceSolutions.com	Computational Science Solutions
Johnson	Calvin	cjohnson@sciences.sdsu.edu	San Diego State University
Johnson	Fred	fcj@acm.org	Science Applications International Corporation
Kerman	Arthur	kerman@mitlns.mit.edu	MIT/UT/ORNL
Khaleel	Moe	moe.khaleel@pnl.gov	Pacific Northwest National Laboratory
Klein	Richard	klein@astro.berkeley.edu	University of California, Berkeley
Kober	Edward	emk@lanl.gov	Los Alamos National Laboratory
Kothe	Doug	Kothe@ornl.gov	Oak Ridge National Laboratory
Lamb	Donald	lamb@oddjob.uchicago.edu	University of Chicago
Landsberg	Sandy	landsberg@ascr.doe.gov	U.S. Department of Energy, Advanced Scientific Computing Research
Larzelere	Alex	alex.larzelere@nuclear.energy.gov	U.S. Department of Energy, Office of Nuclear Energy
Lee	Sander	sander.lee@nnsa.doe.gov	National Nuclear Security Administration
Lee	Stephen	srlee@lanl.gov	Los Alamos National Laboratory
Lee	Tsengdar	tsengdar.j.lee@nasa.gov	National Aeronautics and Space Administration
Liang	Edison	liang@rice.edu	Rice University
Lightstone	Felice	felice@llnl.gov	Lawrence Livermore National Laboratory
Lindsay	Robert	robert.lindsay@science.doe.gov	U.S. Department of Energy, Office of Science
Little	Robert	rcl@lanl.gov	Los Alamos National Laboratory
Livny	Miron	miron@cs.wisc.edu	University of Wisconsin
Maccabe	Barney	maccabeab@ornl.gov	Oak Ridge National Laboratory
Mailhiot	Christian	mailhiot1@llnl.gov	Lawrence Livermore National Laboratory
Marques	Osni	osni.marques@science.doe.gov	U.S. Department of Energy, Office of Science
Matthews	Hope	hope.matthews@pnl.gov	Pacific Northwest National Laboratory
Mattson	Thomas	trmatts@sandia.gov	Sandia National Laboratories
McCoy	Michel	mccoy2@llnl.gov	Lawrence Livermore National Laboratory
Messina	Paul	messina@mcs.anl.gov	Argonne National Laboratory
Modine	Normand	namodin@sandia.gov	Sandia National Laboratories
Moin	Parvis	moin@stanford.edu	Stanford University
Morss	Sue	sue.morss@anl.gov	Argonne National Laboratory
Navratil	Petr	navratil1@llnl.gov	Lawrence Livermore National Laboratory
Nazarewicz	Witold	witek@utk.edu	University of Tennessee
Ng	Esmond	egng@lbl.gov	Lawrence Berkeley National Laboratory

Last Name	First Name	Email	Affiliation
Nichols	Jeffrey	nicholsja@ornl.gov	Oak Ridge National Laboratory
Nikkel	Daniel	dan.nikkel@nnsa.doe.gov	National Nuclear Security Administration
Nowell	Lucy	lucy.nowell@ascr.doe.gov	U.S. Department of Energy, Office of Science
Oehmen	Chris	christopher.oehmen@pnl.gov	Pacific Northwest National Laboratory
Ormand	William	ormand1@llnl.gov	Lawrence Livermore National Laboratory
Pao	Karen	karen.pao@science.doe.gov	U.S. Department of Energy, Advanced Scientific Computing Research
Pederson	Mark	mark.pederson@science.doe.gov	U.S. Department of Energy, Office of Science
Plimpton	Steve	sjplimp@sandia.gov	Sandia National Laboratories
Polansky	Walt	walt.polansky@science.doe.gov	U.S. Department of Energy, Office of Science
Pothen	Alex	apothen@purdue.edu	Purdue University
Quaglioni	Sofia	quaglioni1@llnl.gov	Lawrence Livermore National Laboratory
Redondo	Antonio	redondo@lanl.gov	Los Alamos National Laboratory
Reis	Victor	Victor.reis@nnsa.doe.gov	U.S. Department of Energy
Rintoul	Danny	mdrinto@sandia.gov	Sandia National Laboratories
Roberto	Jim	robertojb@ornl.gov	Oak Ridge National Laboratory
Roche	Ken	kenneth.roche@pnl.gov	Pacific Northwest National Laboratory
Rosner	Robert	r-rosner@uchicago.edu	University of Chicago
Sarrao	John	sarrao@lanl.gov	Los Alamos National Laboratory
Schwegler	Eric	schwegler1@llnl.gov	Lawrence Livermore National Laboratory
Seager	Mark	dancerios1@llnl.gov	Lawrence Livermore National Laboratory
Sefcik	Joseph	sefcik1@llnl.gov	Lawrence Livermore National Laboratory
Spotz	Bill	wfspotz@ascr.doe.gov	U.S. Department of Energy, Advanced Scientific Computing Research
Srivastava	Jaideep	jaideep@srivastava@gmail.com	University of Minnesota
Stan	Marius	mastan@lanl.gov	Los Alamos National Laboratory
Stark	Philip	stark@stat.berkeley.edu	University of California, Berkeley
Stevens	Rick	stevens@anl.gov	Argonne National Laboratory
Stewart	James	jrstewa@sandia.gov	Sandia National Laboratories
Strachan	Alejandro	strachan@purdue.edu	Purdue University
Strayer	Michael	michael.strayer@science.doe.gov	U.S. Department of Energy, Office of Science
Streitz	Fred	streitz1@llnl.gov	Lawrence Livermore National Laboratory
Thompson	Ian	i-thompson@llnl.gov	Lawrence Livermore National Laboratory
Thompson	Aidan	athomps@sandia.gov	Sandia National Laboratories
Truong	Trieu	trieu.truong@nnsa.doe.gov	National Nuclear Security Administration
Turner	John	turnerja@ornl.gov	Oak Ridge National Laboratory
Van Duin	Adri	acv13@psu.edu	Pennsylvania State University
Van Eeckhout	Edward	emvan@lanl.gov	Los Alamos National Laboratory
Vannoy	Lucy	lucy.vannoy@pnl.gov	Pacific Northwest National Laboratory
Vary	James	jvary@iastate.edu	Iowa State University
Vashishta	Priya	priyav@usc.edu	University of Southern California
Voter	Arthur	afv@lanl.gov	Los Alamos National Laboratory
Vranas	Pavlos	vranas1@llnl.gov	Lawrence Livermore National Laboratory
Wade	Douglas	douglas.wade@nnsa.doe.gov	National Nuclear Security Administration
Wagner	Albert	wagner@tcg.anl.gov	Argonne National Laboratory
Webster	Robert	robert.webster@nnsa.doe.gov	National Nuclear Security Administration/Los Alamos National Laboratory
White	Andrew	abw@lanl.gov	Los Alamos National Laboratory

Last Name	First Name	Email	Affiliation
Williams	Timothy	zippy@anl.gov	Argonne National Laboratory
Wills	John	jxw@lanl.gov	Los Alamos National Laboratory
Wiringa	Robert	wiringa@anl.gov	Argonne National Laboratory
Yip	Sid	syip@mit.edu	Massachusetts Institute of Technology
Younes	Walid	younes1@llnl.gov	Lawrence Livermore National Laboratory
Zika	Michael	zika@llnl.gov	Lawrence Livermore National Laboratory
Zimmerman	George	gzimmerman@llnl.gov	Lawrence Livermore National Laboratory

APPENDIX 3: ACRONYMS AND ABBREVIATIONS

ADLB	Asynchronous Dynamic Load Balancing Library
ASC	Advanced Simulation and Computing
ASCR	U.S. Department of Energy Office of Advanced Scientific Computing Research
ASR	adaptive sample refinement
ATD	diabatic time-dependent
BSC	basic science challenges
CC	coupled cluster
CFD	computational fluid dynamics
CO_2	carbon dioxide
CPU	central processing unit
DARHT	Dual Axis Radiographic Hydrodynamic Test Facility
DFT	density functional theory
DHS	U.S. Department of Homeland Security
DMFT	dynamical mean field theory
DNA	deoxyribonucleic acid
DoD	U.S. Department of Defense
DOE	U.S. Department of Energy
EMP	Electro-Magnetic Pulse
EXX/OEP	exact exchange/optimized effective potential methods
FFT	Fast Fourier Transforms
GeV	gigaelectronvolt
GFMC	Green's function Monte Carlo
GHG	greenhouse gas
HANE	High Altitude Nuclear Explosions
HFB	Hartree-Fock-Bogoliubov
HPC	high-performance computing
HTTP	hypertext transfer protocol

ICF	inertial confined fusion
IP	interatomic potentials
KMC	kinetic Monte Carlo
LACM	large-amplitude collective motion
LANSCE	Los Alamos Neutron Science Center
LQCD	lattice quantum chromodynamics
MD	molecular dynamics
MeV	megaelectronvolt
meV	millielectronvolt
NCSM	No-Core Shell Model
NIC	National Ignition Campaign
NIF	National Ignition Facility
NN	two-nucleon
NNN	three-nucleon
NNSA	National Nuclear Security Administration
NRC	U.S. Nuclear Regulatory Commission
NWS	National Weather Service
ODE	ordinary differential equations
PDF	probability density functions
PES	potential energy surface
PFLOP	peta floating point operations per second
pRAD	proton radiography
PRD	priority research directions
QCD	quantum chromodynamics
QMC	quantum Monte Carlo
QMD	quantum molecular dynamics
RAID	redundant array of inexpensive disks
RGM	Resonating Group Method
RMD	reactive molecular dynamics

SCC	stress corrosion cracking
SciDAC	Science Discovery through Advanced Computing
SCP	self-consistent phonon
SNIF-NMR	Site-Specific Natural Isotope Fractionation-Nuclear Magnetic Resonance
ST&E	science, technology, and engineering
TDGCM	time-dependent generator coordinate method
TEP	Trapping of Energetic Particles
UNEDF	Universal Nuclear Energy Density Functional
UQ	uncertainty quantification

U.S. DEPARTMENT OF
ENERGY

Office of Science

Office of Advanced Scientific Computing Research
National Nuclear Security Administration

Prepared for the U.S. Department of Energy under Contract DE-AC05-76RLO1830

Production support provided by Pacific Northwest National Laboratory,
Fundamental & Computational Sciences Directorate

PNNL-19379