Scientific Grand Challenges

FOREFRONT QUESTIONS IN NUCLEAR SCIENCE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE

January 26-28, 2009 • Washington, D.C.

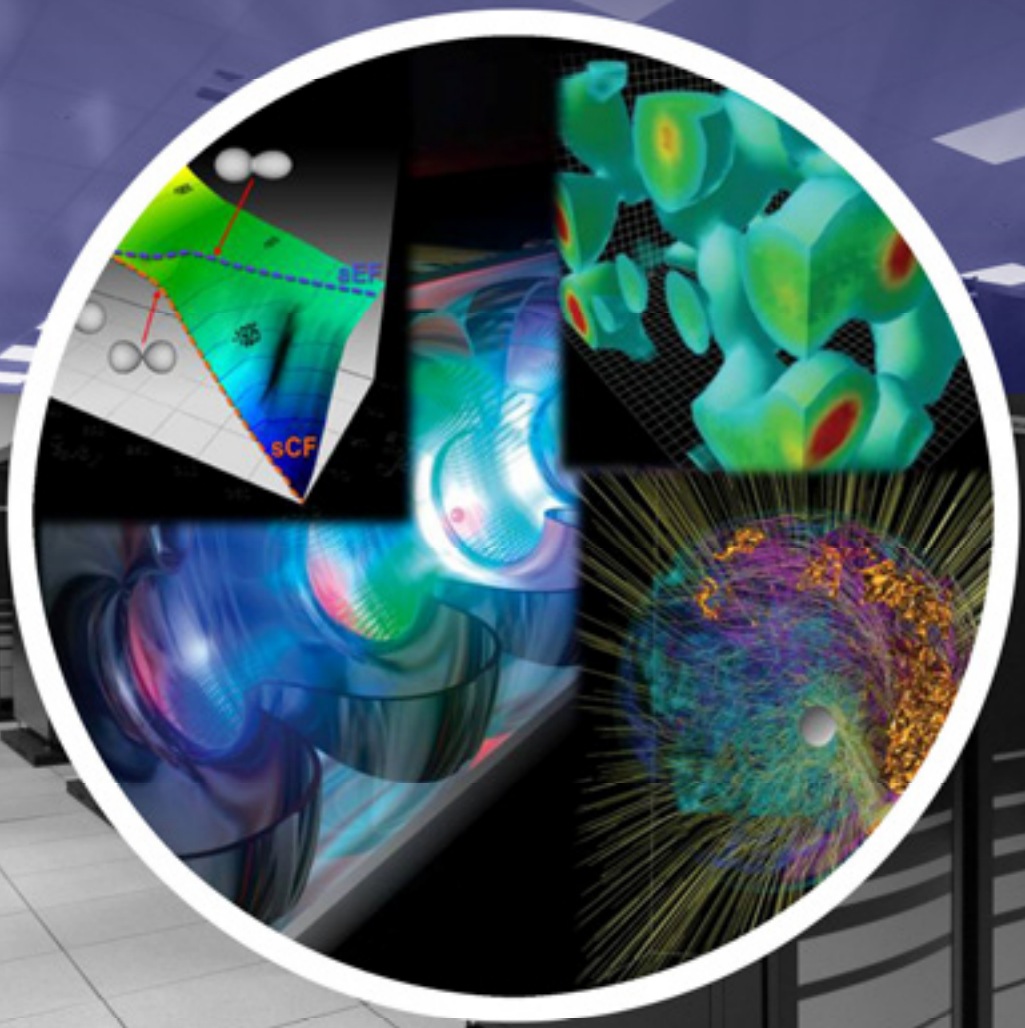

U.S. DEPARTMENT OF ENERGY

Sponsored by the Office of Nuclear Physics and the Office of Advanced Scientific Computing Research

DISCLAIMER

This report was prepared as an account of a workshop sponsored by the U.S. Department of Energy. Neither the United States Government nor any agency thereof, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Copyrights to portions of this report (including graphics) are reserved by original copyright holders or their assignees, and are used by the Government's license and by permission. Requests to use any images must be made to the provider identified in the image credits.

On the cover: Lawrence Berkeley National Laboratory's Cray XT4™ supercomputer. The computer is named Franklin in honor of Benjamin Franklin, one of the first American scientists. The Franklin is at the National Energy Research Scientific Computing Center, which is managed by Lawrence Berkeley National Laboratory for the U.S. Department of Energy's Office of Science. The Franklin's computing power makes it one of the largest and most powerful supercomputers in the world. Future reports in the Scientific Grand Challenges workshop series will feature different Office of Science computers on their covers.

SCIENTIFIC GRAND CHALLENGES: FOREFRONT QUESTIONS IN NUCLEAR SCIENCE AND THE ROLE OF COMPUTING AT THE EXTREME SCALE

Report from the Workshop Held January 26-28, 2009

Sponsored by the U.S. Department of Energy, Office of Nuclear Physics and the Office of Advanced Scientific Computing

Chair, **Glenn R. Young**
Oak Ridge National Laboratory

Co-Chair, **David J. Dean**
Oak Ridge National Laboratory

Co-Chair, **Martin J. Savage**
University of Washington

Workshop Facilitator, **Fred E. Bertrand, Jr.**
Oak Ridge National Laboratory

Co-Lead, *Nuclear Structure and Nuclear Reactions,* **James P. Vary**
Iowa State University

Co-Lead, *Nuclear Structure and Nuclear Reactions,* **Steven C. Pieper**
Argonne National Laboratory

Co-Lead, *Nuclear Astrophysics*, **Anthony Mezzacappa**
Oak Ridge National Laboratory

Co-Lead, *Nuclear Astrophysics*, **George Fuller**
University of California, San Diego

Co-Lead, *Cold Quantum Chromodynamics and Nuclear Forces*, **Thomas Luu**
Lawrence Livermore National Laboratory

Co-Lead, *Cold Quantum Chromodynamics and Nuclear Forces*, **David Richards**
Thomas Jefferson National Accelerator Facility

Co-Lead, *Hot and Dense Quantum Chromodynamics*, **Steffen A. Bass**
Duke University

Co-Lead, *Hot and Dense Quantum Chromodynamics*, **Frithjof Karsch**
Brookhaven National Laboratory

Lead, *Accelerator Physics*, **Robert Ryne**
Lawrence Berkeley National Laboratory

Representative, Office of Nuclear Physics, **Eugene A. Henry**

Representative, Office of Advanced Scientific Computing Research, **Lali Chatterjee**

EXECUTIVE SUMMARY

Nuclear physics is manifest in areas as seemingly disparate as the history of the early universe, the generation of energy in the sun, and the creation of nearly all the elements in stellar furnaces and explosions. A major focus of nuclear physics is the strong force and the atomic nuclei whose binding is a direct result of it, and whose stability underlies that of the atoms and thus molecules forming the familiar matter of all life forms and everyday objects. The strong nuclear force is a complex one, much more so than the gravitational and electromagnetic forces familiar from daily life. Understanding this nuclear force places particularly difficult demands not only on experiment but also on theory and calculation, particularly on computational power needed for these calculations.

Nuclear physics pursues several major areas of research, including the following:

1. Structure of nuclei
2. Nuclear forces and quantum chromodynamics (QCD), the quantum field theory of the strong interaction
3. Fundamental interactions and symmetries that govern the interactions and the quantum states observed
4. Nuclear astrophysics and the synthesis of nuclei in stars and elsewhere in the cosmos
5. Hot QCD for the highest temperatures known and phases of strongly interacting matter
6. Science and design of the accelerators, detectors, and other tools essential for pursuing these scientific questions.

The analysis and interpretation of results, and the development of an encompassing theoretical and intellectual framework depend on significant and focused investments, particularly in large-scale computing facilities and their ancillary capabilities for data visualization, storage, retrieval and transmittal. These are necessary in parallel to the perhaps more familiar investments in state-of-the-art accelerators and detectors to support experimental investigations. Most forefront problems now driving active research depend on access to large-scale computing facilities for tasks ranging from designing the facilities, making predictions based on current theories, and determining the implications for scientists' overall understanding of the experimental results, all with more demanding requirements for precision.

The scale of computing facilities needed for state-of-the-art scientific exploration far outweighs what is possible with desktop- to departmental-scale (few to hundreds of computing cores) resources—instead, requiring access to the (few) large national centers where tens of thousands of computing cores, capable of delivering sustained performance of hundreds of teraflops, are available. Many problems, coming from all areas of nuclear physics, still require that significant approximations be made before such powerful centers can process the problem, or cannot even be attempted at present centers as they require computing timescales on the order of years. Extrapolating to the computing power needed to remove such approximations indicates a major shift in the ability to address such forefront questions that will occur with computing resources that are a few orders of magnitude more capable than current state-of-the-art machines. It can be projected that in a number of areas of active research, descriptions become possible that include all presently known phenomena and address the science involved at the multitude of physics scales that are known to be important. This report provides a focused consideration of the nuclear

physics that could be addressed at such a facility and notes proactive investments in computing facilities that would advance state-of-the-art technology.

The shift in architecture away from single-core to multicore computer architectures over the past two decades means that efforts to harness and make best use of the computing power offered need not only involve scientists and applied mathematicians to develop the best numerical approaches to solving computational problems, but necessarily also involves collaboration with computer scientists to develop programming approaches and languages that make best use of the massively parallel architectures. Thus, any exploration of the possibilities for computation in nuclear physics at scales that exceed those of today necessarily involve joint efforts among scientists, applied mathematicians, and computer scientists.

This report is an account of the deliberations and conclusions of the workshop on "Forefront Questions in Nuclear Science and the Role of High Performance Computing" held January 26-28, 2009, co-sponsored by the U.S. Department of Energy (DOE) Office of Nuclear Physics and the DOE Office of Advanced Scientific Computing Research. Representatives from the national and international nuclear physics communities, as well as from the high-performance computing community, participated. The purpose of this workshop was to 1) identify forefront scientific challenges in nuclear physics and then determine which—if any—of these could be aided by high-performance computing at the extreme scale; 2) establish how and why new high-performance computing capabilities could address issues at the frontiers of nuclear science; 3) provide nuclear physicists the opportunity to influence the development of high-performance computing; and 4) provide the nuclear physics community with plans for development of future high-performance computing capability by the DOE Office of Advanced Scientific Computing Research.

Technical panel discussions focused on five major areas where extreme scale computing is most relevant to nuclear science:

- nuclear forces and cold QCD
- nuclear structure and nuclear reactions
- nuclear astrophysics
- hot and dense QCD
- accelerator physics.

This report gives a description of key science issues and opportunities for these five major science areas. The current status is exhibited with examples drawn from active research that uses current large computing centers. The report then provides priority research directions (PRDs), including scientific goals that can be achieved at each step of a series of measured increases in computing power above that presently available, culminating in a factor of one-thousand-fold increase in capability. A discussion of crosscutting issues is provided, with opportunities featured for collaboration with scientists particularly in the areas of applied mathematics and computer science. The report concludes with a conclusions and recommendations section.

Priority Research Directions

The PRDs that the five sub-panels identified are summarized in the following paragraphs.

Sub-Panel 1 – Nuclear Forces and Cold Quantum Chromodynamics

QCD is the underlying quantum field theory describing the strong interactions. When combined with the electroweak interactions, it is responsible for the interactions and properties of the nucleons and all nuclei. Lattice QCD is the only known technique which can be used to solve QCD and involves a direct numerical evaluation of the path-integral of QCD using large-scale computers. This sub-panel identified four PRDs.

- The first is **the spectrum of QCD**, which involves the calculation of the spectrum and properties of excited states of mesons, in particular those with explicit gluonic degrees of freedom.
- The second is **how QCD makes a proton**, which involves the calculation of the contribution of gluons to nucleon structure, which will refine scientists' knowledge of the origin of mass and spin in the proton.
- The third is **from QCD to nuclei**, which involves the calculation of the three-nucleon interaction with lattice QCD, which will provide interactions between nucleons that must be known, yet are inaccessible to experimentation; this inaccessibility currently limits nuclear structure/reaction and astrophysics calculations.
- The fourth is **fundamental symmetries**, which involves the calculation of both parity-violating nuclear forces and time-reversal-violating observables resulting from the weak- and strong-interactions with lattice QCD.

Sub-Panel 2 – Nuclear Structure and Nuclear Reactions

The theory of the atomic nucleus focuses on predicting and explaining extensive classes of phenomena that occur in nuclei. These phenomena are important in the birth of the universe, in astrophysical settings, in energy generation via fission and fusion, and in industrial and medical applications via use of stable isotopes and radioisotopes. This sub-panel identified four PRDs.

- The first is the **physics of extreme neutron-rich nuclei and matter**, which includes computing properties of nuclei that determine the r-process nucleo-synthesis path in stars and the nucleonic matter in neutron star cores and crusts.
- The second is the **microscopic description of nuclear fission**, which involves solving a problem of both basic science interest and of great practical importance—computing the paths to fission and its products.
- The third is **nuclei as neutrino physics laboratories**, which involves computing properties of nuclei used in double-beta decay experiments and neutrino-nucleus cross sections for modeling supernova explosions and for the exploration of the properties of neutrinos.
- The fourth is the **reactions that made us**, which involves computing the triple-alpha process that produces ^{12}C, the nucleus at the core of organic chemistry and thus life forms, and $^{12}C(\alpha,\gamma)^{16}O$, the element that is key to both water and the reactions that power humans and our present society.

Sub-Panel 3 – Nuclear Astrophysics

The study of astrophysics and cosmology helps scientists understand the universe, how it came to be, and our place in it. Given ever more precise astrophysical models, the universe becomes a far more powerful laboratory for neutrino physics beyond the standard model, astrophysics, and nuclear physics at high density and neutron content. Laboratory experiments in nuclear and neutrino physics in turn can influence understanding of astrophysics. This sub-panel identified three PRDs.

- The first is **the sun and other stars**, which involves both developing the first three-dimensional models of the sun and other stars that capture the turbulence of stellar interiors, and determining their ramifications for stellar evolution, nucleosynthesis, and neutrino physics.
- The second is **stellar explosions and their remnants: thermonuclear supernovae**, which involves the development of high-fidelity models of thermonuclear supernovae (the standard candles for studying the behavior of dark energy with redshift) with complete treatments of turbulent nuclear burning, of the transition from deflagration to detonation, and of radiation transfer to determine their explosion mechanism and element synthesis.
- The third is **stellar explosions and their remnants: core-collapse supernovae**, which involves the development of precision three-dimensional models of core-collapse supernovae to ascertain their explosion mechanism; to predict their element synthesis in both explosive and r-process channels; to predict their neutrino and gravitational wave signatures; and to study their remnants (e.g., neutron stars), which then leads to understanding neutrino and nuclear physics not accessible in, and/or complementing, terrestrial experiments.

Sub-Panel 4 – Hot and Dense Quantum Chromodynamics

The recent discovery of deconfined strongly interacting matter, colloquially called quark-gluon plasma, is a significant opportunity to extend current understanding of nuclear matter and elementary particles at high temperature and density. This matter shows evidence of behaving as a fluid with a remarkably small ratio of shear viscosity to entropy, possibly lower than any other known substance. There is evidence from spectral shapes, spectra of heavy quarks, and the elliptical asymmetry of emitted particles, that equilibration occurs on a remarkably small timescale. This sub-panel identified four PRDs.

- The first is the **precision calculation of bulk thermodynamics** for strongly interacting matter.
- The second is the **QCD phase structure at nonzero net baryon number density**, with a goal to calculate the behavior of QCD at finite-temperature and density to map the QCD phase-diagram, and determine the existence and location (if it exists) of the thus far elusive critical point.
- The third is the **transport coefficients of QCD and spectral functions of hadrons in medium**, which involves calculating the transport properties of QCD at high temperatures and nonzero densities, including the screening mechanisms that lead to modifications of hadron properties and the destruction of light and heavy quark hadronic bound states.
- The fourth is the **equilibration challenge: from the color glass condensate to the quark gluon plasma**, which involves using the outputs of the first three PRDs in the complex three-dimensional modeling of the expansion and cooling of dense matter created in relativistic heavy ion collisions.

Sub-Panel 5 – Accelerator Physics

Accelerators are essential elements of the nuclear physics program, with the majority of experiments requiring access to advanced accelerators for their execution. The recently announced Facility for Rare Isotope Beam physics will have a state-of-the-art heavy-ion linac and separation system to obtain rarely produced nuclei at production levels of 10^{-17} or less of the primary beam for key experiments in nuclear structure and astrophysics. Cold and hot QCD physics programs both require knowledge of the structure of nuclei that can only be accessed by designing a new type of colliding-beams accelerator that can collide electrons and heavy atomic nuclei at luminosities that are orders of magnitude beyond state-of-the-art technology. Design tools used to realize these machines must address questions ranging from detailed long-term stability of particle orbits and beam emittances to detailed modeling of electromagnetic structures used to support the accelerating fields. This sub-panel identified four PRDs.

- The first is to **maximize the production efficiency and beam purity for rare isotope beams for nuclear physics experiments**, which requires methods to model beam separators to remove beam fragmentation products to the 10^{-17} level.
- The second is to **develop an optimal design for an electron-ion collider**; such a machine will operate at unprecedented luminosities and will require advanced modeling of the beam-cooling channels and related devices to control beam halo and disruption.
- The third is **design optimization of complex electromagnetic structures for nuclear physics accelerator facilities**. New accelerators tend to be large and represent significant capital investment; therefore, a means of modeling the electromagnetic response of accelerator cavity and magnetic structures incorporating multi-physics is needed.
- The fourth is **advanced methods of accelerator simulation**. New concept accelerators require much exploration of parameter space to determine if the basic idea will produce an interesting machine.

TABLE OF CONTENTS

EXECUTIVE SUMMARY V
TABLE OF CONTENTS XI
INTRODUCTION 1
PANEL REPORTS 5
COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES 7
Introduction and Current Status 7
Priority Research Directions 15
Applied Mathematics and Computational Issues 29
Enhanced Synergism With Other Subfields of Nuclear Physics and the Broader Physics Community 32
Conclusions 33
NUCLEAR STRUCTURE AND NUCLEAR REACTIONS 35
Introduction and Current Status 35
Priority Research Directions 36
Conclusions 49
NUCLEAR ASTROPHYSICS 51
Introduction and Current Status 51
Priority Research Directions 54
Crosscutting Research Direction: Light Curves and Spectra from Thermonuclear and Core-Collapse Supernovae 68
Tools for Other Astrophysical Systems 71
Building a Common Vocabulary in Applied Mathematics and Computer Science 73
Conclusions 74
HOT AND DENSE QUANTUM CHROMODYNAMICS 77
Introduction and Current Status 77
Priority Research Directions 88
Applied Mathematics and Computational Issues 98
Enhanced Synergism with Other Subfields of Nuclear Physics and the Broader Physics Community 99
Conclusions 100
ACCELERATOR PHYSICS 103
Introduction and Current Status 103
Priority Research Directions 119
Applied Mathematics and Computational Issues 130
Conclusions 131

CROSSCUTTING CHALLENGES 133
CROSSCUTTING CHALLENGES 135
PRIORITY RESEARCH DIRECTIONS 139
COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES 141
Spectrum of Quantum Chromodynamics 141
How Quantum Chromodynamics Makes a Proton 143
From Quantum Chromodynamics to Nuclei 146
Fundamental Symmetries 149
NUCLEAR STRUCTURE AND NUCLEAR REACTIONS 153
Ab Initio Calculations of Light Nuclei and Their Reactions 153
Reactions That Made Us: Triple-Alpha Process and $^{12}C(\alpha,\gamma)^{16}O$ 153
Weak Nuclear Structure—Nuclei as Laboratories for Neutrino Physics 156
Microscopic Description of Nuclear Fission 159
Physics of Extreme Neutron-Rich Nuclei and Matter 162
NUCLEAR ASTROPHYSICS 165
Sun and Other Stars 165
Stellar Explosions and Their Remnants: Thermonuclear Supernovae 169
Stellar Explosions and Their Remnants: Core-Collapse Supernovae 172
HOT AND DENSE QUANTUM CHROMODYNAMICS 177
Precision Calculation of Bulk Thermodynamics 177
Quantum Chromodynamics Phase Structure at Nonzero Net Baryon Number Density 180
Transport Coefficients of Quantum Chromodynamics and Spectral Functions of Hadrons in Medium 182
Equilibration Challenge: From the Color Glass Condensate to the Quark- Gluon Plasma 185
ACCELERATOR PHYSICS 189
Maximize Production Efficiency, Variety, and Purity of Rare Isotope Beams for Nuclear Physics Experiments 189
Develop Optimal Design for an Electron-Ion Collider 192
Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities 194
Advanced Methods and Applications of Accelerator Simulations for Nuclear Physics Facilities 197
CONCLUSIONS AND RECOMMENDATIONS 201
Conclusions 201
Recommendations 205

REFERENCES 207
APPENDIX 1: WORKSHOP AGENDA ..... APPENDIX 1-1
APPENDIX 2: WORKSHOP PARTICIPANTS ..... APPENDIX 2-1
APPENDIX 3: ACRONYMS AND ABBREVIATIONS ..... APPENDIX 3-1

INTRODUCTION

The strong nuclear forces—or strong interactions—give rise to a wealth of phenomena that govern the formation of atomic nuclei in the cosmos. These forces are responsible for the binding and stability of atomic nuclei and thus the molecules forming the familiar matter of everyday experience, including life itself. In contrast to the gravitational and electromagnetic forces that are familiar to all humans, the strong nuclear forces are particularly complex, and their study defines the core of nuclear physics. Developing a complete understanding of their implications for the universe and society places extreme demands not only on experiment, but also on theory. A broad range of investigations are underway to improve and expand current understanding of these strong nuclear forces. This includes understanding and quantifying how they result in the observed properties and interactions of nuclei and their constituents, and how they are involved in areas as seemingly disparate as the history of the early universe, the generation of energy in the sun and in terrestrial nuclear reactors, and the creation of nearly all the elements in stellar furnaces and explosions.

Nuclear physics pursues several major areas of research to examine these many aspects, including research into the structure of nuclei; into nuclear forces and quantum chromodynamics (QCD) (the quantum field theory of the strong interaction); into the fundamental interactions and the symmetries that govern these interactions and the observed quantum states; into nuclear astrophysics and the synthesis of nuclei in stars and elsewhere in the cosmos; into hot QCD and phases of strongly interacting matter; and in the science and design of the accelerators, detectors, and other tools essential for pursuing these scientific questions. The active areas of forefront research are described in the *The Frontiers of Nuclear Science, A Long Range Plan* (DOE 2007), including progress towards various goals set for that research.

Significant and focused investments are needed to support much of the planned research that is outlined above. Much of the experimental investigation requires use of state-of-the-art accelerators and detection devices, as outlined in the 2007 Nuclear Science Long Range Plan (DOE 2007). The analysis and interpretation of results, and the development of an encompassing theoretical and intellectual framework, also depend in a major way on such focused investments, particularly on large-scale computing facilities and their ancillary capabilities for data visualization, storage, retrieval and transmittal. An examination of most active lines of current research shows that many of the forefront problems now driving active research inquiries depend on access to large-scale computing facilities for tasks ranging from designing the facilities, making predictions based on current theories, and working out the implications for overall understanding of the experimental results.

The scale of computing facilities needed for forefront research in many cases far outstrips what is possible with current desktop-scale (one to four computing cores) and current university departmental-scale (dozens to hundreds of computing cores) or even most institutional-scale (thousands of computing cores) resources. Instead, this research requires access to the few large national centers where tens of thousands of computing cores—capable of delivering sustained performance of multiple teraflops and recently petaflops—are available. However, it remains the situation that the resources currently available at computational centers are insufficient to solve key problems in all areas of nuclear physics without either controlled or uncontrolled approximations. Extrapolating to the computing power needed to remove uncontrolled approximations and quantify controlled approximations indicates that a major shift in the ability to address forefront questions in nuclear physics will occur with computing capability and capacity that are a few orders of magnitude larger than what is available today. It can be projected that, in a

number of areas of active research, descriptions become possible that include all presently known phenomena and address the science involved at the multitude of physics scales known to be important. A focused consideration of nuclear physics that requires—or could be addressed with—such facilities is thus timely.

The evolution of large computer capability has led to large machines capable of sustaining petaflops, with further growth definitely possible with known architectures. This represents a remarkable increase in top speed in just over two decades. The Cray-2 attained one gigaflop in 1986, the Intel® ASC Red attained one teraflop in 1997, the IBM® Roadrunner attained one petaflop in 2008, and was soon followed by the Cray XT5™.[1,2,3,4] The evolution of processor speeds has also been dramatic, although less so than the overall performance of very large machines, moving from the 8-bit processors used in early personal computers such as the Apple® II, which ran at 1 megahertz, to modern 64-bit microprocessors operating at 3.5 gigahertz.[5] Note that processor speeds have increased little in the past 6 years, as technical limits and issues of heat removal from the silicon chips employed have come to dominate technical development; indeed, much engineering development is presently focused on increasing the number of computing cores located on one substrate. The greater scale up in aggregate performance achieved by the largest machines has been achieved by parallel processing where hundreds, thousands, and now hundreds of thousands of computing cores are focused on one problem: the five largest multiprocessor machines on the TOP500 list have over 100,000 computing cores. In addition, future machines will include accelerator technology, effectively bringing parallelism into the millions of computing cores.

This shift in architecture away from single processors over the past two decades means that efforts to harness and make best use of the available computing power not only involves scientists and applied mathematicians to develop the best numerical approaches to solving computational problems, but necessarily involves collaboration with computer scientists to develop programming approaches and languages that make best use of the massively parallel architectures. Thus, any exploration of the possibilities for computation in nuclear physics at scales larger than today must involve joint efforts among scientists, applied mathematicians, and computer scientists.

This workshop report is one of a series resulting from the Scientific Grand Challenges Workshops hosted by the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) in partnership with other DOE 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, interim director of science at the Argonne Leadership Computing Facility, is overseeing the workshop series.

The workshop, “Forefront Questions in Nuclear Science and the Role of High Performance Computing,” was held January 26-28, 2009, in Gaithersburg, Maryland, to discuss leading scientific problems in nuclear physics with a further focus on those problems that require extreme scale scientific computing capabilities. This collaborative workshop, part of the Scientific Grand Challenges Workshops series, was

[1] Intel is a trademark of Intel Corporation.
[2] IBM is a trademark of International Business Machines Corporation.
[3] Cray XT5 is a trademark of Cray Inc.
[4] TOP500 is a trademark of TOP500. Further information is available at TOP500® Supercomputer Sites at http://www.top500.org.
[5] Apple is a trademark of Apple Inc.

co-sponsored by the DOE Office of Nuclear Physics and DOE ASCR. A copy of the charge letter requesting the workshop is provided in Young (2008).

Details of the workshop program, the overall workshop charge letter from DOE (Young 2008), and the plenary presentations are provided on the Extreme Scale Computing Workshop - Nuclear Physics website at http://extremecomputing.labworks.org/nuclearphysics/index.stm. Workshops in this series have been completed in climate science, high energy physics, nuclear physics, fusion energy sciences, nuclear energy, biology, and basic energy sciences. A workshop on national security will be conducted in October 2009.

The purpose of this workshop was to 1) identify forefront scientific challenges in nuclear physics and then determine which—if any—of these could be aided by high performance computing at the extreme scale; 2) establish how and why new high performance computing capabilities could address issues at the frontiers of nuclear science; 3) provide nuclear physicists the opportunity to influence the development of high performance computing; and 4) provide the nuclear physics community with plans for development of future high performance computing capability by DOE ASCR. The workshop provided a forum for the exchange of ideas from multiple scientific disciplines, including nuclear physicists, computer scientists, and applied mathematicians. These participants provided the interdisciplinary expertise required to identify and address challenges in nuclear science and high performance computing, with an emphasis on the use of extreme scale computing for nuclear science research, advances, and discoveries. The discussions included science issues, resource issues, the needed discipline skills to design and/or adapt code to use large computing centers, and the practical aspects of using very large centers and planning for future centers.

One hundred and fourteen participants registered for this workshop, of which 109 were in attendance. Workshop participants included physicists representing multiple areas of investigation in nuclear physics, computer science, and applied mathematics. Participants represented 27 national and foreign universities, 7 DOE national laboratories, 6 national and foreign corporations, and 2 federal agencies. Also present as observers were several DOE Headquarters program managers. The workshop agenda is provided in Appendix 1, followed by a list of attendees and their respective institutions in Appendix 2.

Technical panel discussions focused on five major areas where extreme scale computing is most relevant to nuclear science:

- nuclear forces and cold QCD
- nuclear structure and nuclear reactions
- nuclear astrophysics
- hot and dense QCD
- accelerator physics.

All groups first focused on the key scientific questions in their areas with reference to the 2007 Nuclear Science Long Range Plan (DOE 2007), and the more immediate scientific deliverables planned for the field (milestones). Forefront problems that would benefit from the application of large-scale computing resources were then examined and the relevant scale of current and needed work determined. A common goal of these five technical panel discussions was to define future scientific challenges that require the development of extreme scale computing to be addressed.

The structure of this report is as follows. A description of key science issues and opportunities for advancement are provided in separate sections for each of the five major areas listed above. The current status is exhibited with examples from active research efforts that draw upon large computing centers for the needed computing, visualization, and data management support. This is followed by the priority research directions for each of the five areas. For each priority research direction, there is a discussion of the scientific goals that can be achieved with a series of measured increases in computing power above that presently available, with a factor of one-thousand-fold increase in capability representing the highest level considered. Crosscutting issues are then presented, with opportunities featured for collaboration with scientists particularly in the areas of applied mathematics and computer science. Conclusions and recommendations from this report are provided, followed by references to published literature cited in the report. The workshop agenda is provided in Appendix 1, followed by a list of attendees and their respective institutions in Appendix 2. Appendix 3 provides a list of acronyms and abbreviations used throughout this report.

PANEL REPORTS

COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES

NUCLEAR STRUCTURE AND NUCLEAR REACTIONS

NUCLEAR ASTROPHYSICS

HOT AND DENSE QUANTUM CHROMODYNAMICS

ACCELERATOR PHYSICS

COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES

Co-Leads: Thomas Luu, Lawrence Livermore National Laboratory
David Richards, Thomas Jefferson National Accelerator Facility (Jefferson Laboratory)

INTRODUCTION AND CURRENT STATUS

The strong interaction, one of the fundamental forces of nature, is responsible for a diverse range of physical phenomena, including binding gluons and the lightest quarks into pions, protons, and neutrons. Through its manifestation as the strong nuclear force, it then binds neutrons and protons together to form the elements of the periodic table. The strong nuclear force plays a central role in the burning of light nuclei, or fusion, which occurs within the sun's core. Within certain heavy nuclei—such as the actinides—the strong nuclear force, in a delicate balance with the electromagnetic interaction, is responsible for the process of fission where nuclei split into smaller constituents, a process exploited in nuclear power stations to produce carbon-free energy. In core-collapse supernovae, the strong nuclear force is essential for producing the shockwave that ejects the star's mantle, providing an environment for the nucleosynthesis of elements above iron. In all the subfields of nuclear physics represented in this report, the strong nuclear force plays a central role.

The strong nuclear force originates from an underlying quantum field theory known as quantum chromodynamics (QCD). This theory governs the interactions of quarks and gluons that are basic constituents of the observable matter in our surrounding environment. QCD has been thoroughly tested by experiments at high energies, giving scientists insight into nature's workings over distances that are smaller than the size of nucleons (the term used for both protons and neutrons). However, at low energies or larger distances, the theory becomes formidable and efforts to theoretically determine fundamental nuclear physics phenomena directly from QCD have met with less success. A long-standing effort of the U.S. Department of Energy's (DOE) Nuclear Physics program is understanding how QCD in this low-energy regime—dubbed cold QCD—manifests itself into observed nuclear phenomena, and further, how scientists can use QCD to make reliable predictions for processes that cannot be experimentally accessed.

The key issue is the phenomenon of confinement: quarks, through their interaction with gluons, are never found in isolation at these low energies; they come as quark/antiquark pairs known as mesons, or as triplets of quarks known as baryons. The phenomenon of confinement is a well established outcome of QCD, but the exact nature of how confinement occurs and all of its consequences are still a mystery. Why does QCD give us the particular mesons that we observe? In baryons, such as neutrons or protons, how exactly do the quarks and gluons behave? Answers to these and similar questions will give scientists greater insight into the mechanism of confinement and its consequences. Similar questions remain unanswered about the exact nature of the strong nuclear force and its lineage from QCD. For example, what is the nature and origin of the nuclear spin-orbit interaction? How does the three-neutron force emerge from QCD? Even partial answers to these questions will help illuminate long-standing issues in nuclear physics and ultimately provide scientists with a deeper understanding of how protons and neutrons combine to form nuclei, and therefore the elements germane to everyday life.

What is Quantum Chromodynamics and Lattice Chromodynamics?

At distance scales much less than 10^{-15} m, matter is found to be composed of quarks and gluons. There are six types of quarks, or "flavors": up, down, strange, charm, bottom, and top quarks. They carry fractional electric charges: for example, an up quark carries negative two-thirds of the charge of the electron, whereas a down quark carries one-third of the charge of the electron. Quarks also carry another type of charge that physicists quizzically call "color": any quark will have either red, blue, or green color charge. The theory that governs the interaction of quarks is aptly called quantum chromodynamics (QCD). Within this theory, gluons are massless particles that mediate forces between quarks, as well as between themselves. Gluons carry no electric charge but do carry color charges. At very short-distances, the interaction strength between quarks and gluons becomes small, a feature of QCD called "Asymptotic Freedom" (Gross and Wilczek 1973; Politzer 1973), while at larger distances the interaction between quarks and gluons becomes very strong, confining them into "colorless" objects called hadrons. The 2004 Nobel Prize in physics was awarded to David J. Gross, H. David Politzer, and Frank Wilczek for their discovery of asymptotic freedom in QCD, a non-Abelian gauge theory.

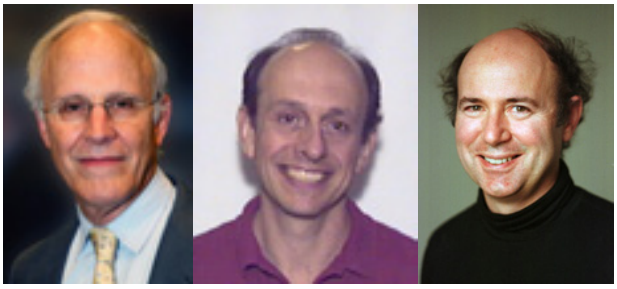

From left to right: David J. Gross, H. David Politzer, and Frank Wilczek. David G. Gross (Kavli Institute, University of California) image courtesy of Los Alamos National Laboratory; H. David Politzer (California Institute of Technology) image courtesy of California Institute of Technology; and Frank Wilczek (Massachusetts Institute Technology professor) image courtesy of Donna Coveney (Massachusetts Institute of Technology).

Because of the peculiar way that gluons interact with each other at low energies, or at distances comparable to or greater than 10^{-15} m, QCD has so far proven impossible to solve using pencil and paper calculations. Lattice QCD is a formulation of QCD in which space and time are discretized on a lattice in such a way that makes the theory amenable to numerical calculations. A lattice QCD calculation is performed using a Monte Carlo method in which samplings of the vacuum are generated with a distribution prescribed by QCD, and physical observables are then measured on these samplings. The greater the number of measurements, the smaller the statistical uncertainty in the calculation. With sufficient computational resources, and a careful account of the discretization artifacts imposed by the lattice, calculations that are not currently possible with pencil and paper are now numerically feasible.

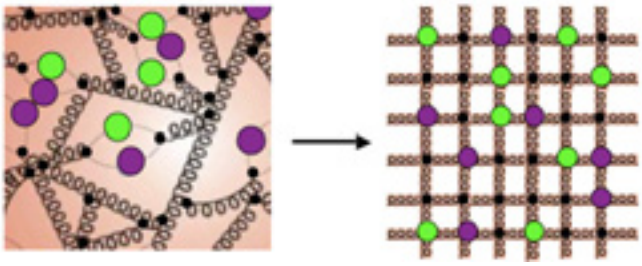

QCD is the quantum field theory that describes the interactions of quarks (depicted as colored circles), and gluons (wiggly lines). Gluons also interact with themselves, as shown in the left cartoon. For numerical calculations of QCD, the theory is formulated on a space-time lattice, as depicted in the right cartoon. Bottom images courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

The breadth of issues outlined in previous paragraphs does not represent a lack of progress in understanding nature in these environments. Rather, computational resources have now reached a point where answers to many of these questions are considered within the realm of possibility. Substantial progress has been made in using available computational resources to investigate these questions. As computational research in these areas matures, scientists can expect further progress in using the increasingly available computational resources. Extreme computing will usher in a new era in computational physics for zero-temperature QCD calculations, and will afford the opportunity of answering many of these outstanding questions in nuclear physics. However, new challenges must be met before such advances in computational resources can be used to their full potential.

Lattice QCD (LQCD) currently provides the only *ab initio* method for performing QCD calculations in the low-energy regime, and for acquiring a quantitative description of the physics of hadrons and nuclear forces. With LQCD calculations at the extreme scale, many hadronic observables will be calculated with precision, and the role of the gluons will be revealed. Aspects of hadronic physics not accessible to experiments will be readily calculated, allowing for many of today's nuclear physics questions to be answered. Many of these calculations will have a direct impact beyond nuclear physics, such as in high-energy physics and in particular, our understanding of the standard model of particle physics.

Effective field theories (EFTs) will play an important role leading up to and beyond the extreme computing era. EFTs represent the natural bridge that connects LQCD to other subfields of nuclear physics. In the extreme computing era, EFTs will be readily constrained by LQCD calculations and used in nuclear structure and nuclear reaction calculations of light nuclei, helping scientists understand how nuclei emerge from QCD. In this manner, EFTs extend the range of applicability of LQCD; processes not directly calculable within LQCD can be investigated through the use of EFTs that are constrained by LQCD. An example of this is using lattice-based calculations of EFTs (LEFTs) to calculate the properties of infinite neutron matter. These calculations will help scientists understand and calculate the properties of neutron star crusts. The impact of extreme computing in this subfield of nuclear physics will be felt by the broader nuclear physics communities, such as those of nuclear structure and reactions, nuclear astrophysics, and nuclear matter under extreme conditions. *The impact of extreme computing on nuclear forces and hadronic physics cannot be overstated.*

The remainder of this panel report is organized as follows. The current status of the field is described, and several salient recent calculations are noted that demonstrate the potential of LQCD for describing the physics of hadrons and nuclei. Following this discussion, four priority research directions (PRDs) are identified that afford the opportunity for extreme computing to transform the current understanding of nuclear physics. For each PRD, the authors describe the computational and scientific challenges that must be met, the expected outcome, and the impact of the outcome on the current understanding of nuclear physics, and other areas of physics. For each PRD, the calculations facilitated through extreme computing are transformational culminations of a long-term research program, and the authors of this report therefore identify a series of intermediate milestone calculations leading to the extreme computing era.

Current Status

There is growing evidence of a synergistic collaboration between various subfields of nuclear physics made possible through the increase in computational resources available to the nuclear physics community. This is especially true in the area of nuclear forces and cold QCD. Based on the

accomplishments listed in this report, it is evident many of these results will have an impact and relevance to other areas of physics. Calculations of the interaction parameters of mesonic systems, for example, will in the future be used in femtoscopy studies of relativistic heavy ion experiments. LQCD calculations of exotic nuclear systems not accessible to experiment will be used in nuclear astrophysics simulations of supernovae. Calculations of moments of structure functions constrain the parton distributions needed to interpret the results of high-energy physics experiments such as those at the Large Hadron Collider in Europe. Recent accomplishments show the overlap of this subfield of nuclear physics with other areas of physics is significant and growing.

Accomplishments over the last 5 years have established the methodology and the groundwork for lattice investigations of hadron structure, spectroscopy, and hadron scattering calculations. An impressive menu of computations has been performed at decreasingly smaller values of the light-quark masses, with present day computational-resource limitations restricting the mass of the pion in precision calculations to be in excess of approximately 300 MeV. In the field of hadron structure, form factors, the lowest three moments of quark, spin, and transversity distributions, and generalized form factors corresponding to the lowest three moments of generalized parton distributions have been calculated (Göckeler et al. 2004; Hägler et al. 2003, 2004). Notable achievements include separating the contributions of the spin and orbital angular momentum of quarks to the nucleon spin (Figure 1), and the observation of the strong dependence of the transverse size of the nucleon on the longitudinal momentum-fraction of the struck quark (Hägler et al. 2008). The transition form factors between the nucleon and $\Delta(1232)$ have been calculated to explore the role of hadronic deformation (Alexandrou et al. 2004, 2005).

In spectroscopy, techniques to calculate nucleon-extended sources within the appropriate representation of the symmetry group of the underlying space-time lattice (the hypercubic group) have been developed (Basak et al. 2005a, 2005b) and used to calculate ground states and excited states (Bulava et al. 2009). Similar techniques have been used to calculate the radiative decay probabilities of excited states of the nucleon such as the $\Delta(1232)$ (Alexandrou et al. 2004, 2005, 2008) and P_{11} (Lin et al. 2008) resonances, as well as the quark distribution amplitudes of the negative-parity partner of the nucleon (Göckeler et al. 2008).

In the meson sector, the charmonium resonance spectrum (the spectrum of states composed of a charm quark and a charm antiquark) has been determined (Chen 2001; Okamoto et al. 2002; Ehmann and Bali 2007; Dudek et al. 2008). Certain electromagnetic properties of charmonium, such as the rate of one- and two-photon emission from their excited states, have also been calculated (Dudek and Edwards 2006; Dudek et al. 2006); see Figure 2. A revolution in the ability to understand the interaction between hadrons, such as neutrons and protons, was the realization that current and future LQCD computations could be used to calculate nuclear interactions without relying on empirical data. LQCD calculations have enabled precise computations of a number of meson-meson interaction parameters (Beane et al. 2008a), demonstrating that rigorous first-principles calculations of the interactions that bind nuclei can be obtained from LQCD. Exploratory calculations are now underway on meson-baryon, baryon-baryon (Beane et al. 2006a), and more generally, multi-hadron systems.

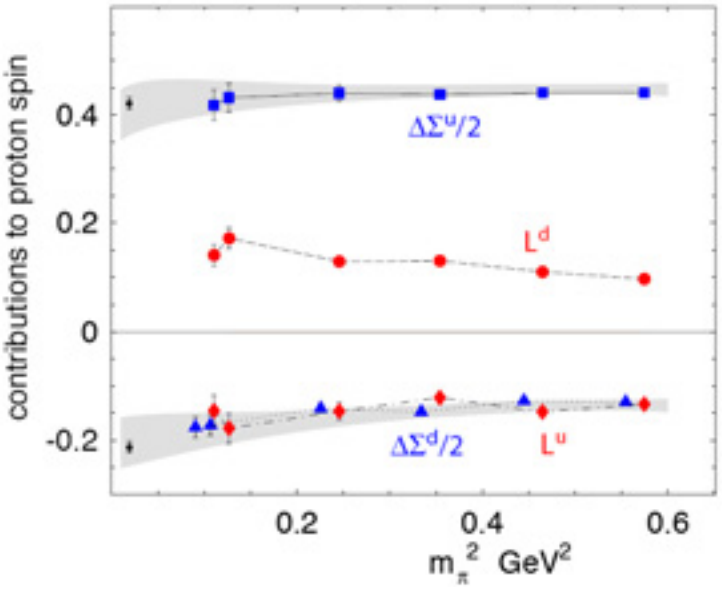

Figure 1. The spin and orbital angular momentum carried by the u and the d quarks in a nucleon as a function of the square of the pion mass, as calculated with lattice quantum chromodynamics. The black points and their uncertainties represent the present experimental determination of their spin contributions. Image courtesy of LHP Collaboration.

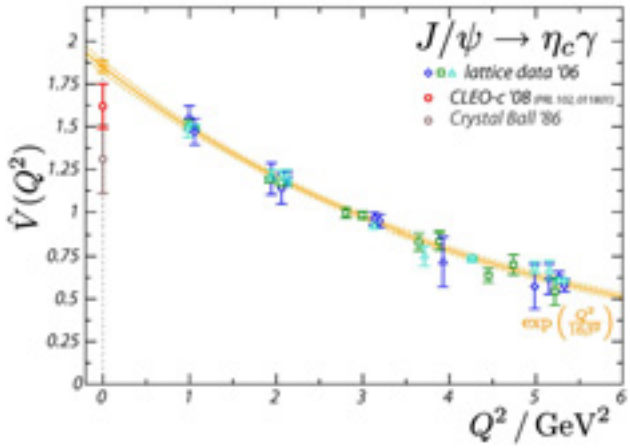

Figure 2. Results of a lattice quantum chromodynamics calculation of the form factor dictating J/ψ to $\eta_c + \gamma^{(*)}$, together with the extrapolation to the case of a real photon (Dudek et al. 2006). Circles denote values of the photo-coupling obtained from experimental measurements of the radiative decay process. Image courtesy of Jozef Dudek (Thomas Jefferson National Accelerator Laboratory).

LQCD calculations are already being used to constrain certain EFTs that are relevant for describing the interactions between mesons (Beane et al. 2008a, 2008b). The EFTs will, in turn, be used to calculate nuclear processes not calculable with LQCD. In the near future, LQCD calculations will be used to constrain the EFTs that are used in nuclear structure calculations of light nuclei, such as those being performed with the no-core shell model (NCSM) (Navrátil and Ormand 2002) and coupled-cluster (CC) (Hagen et al. 2007b) algorithms. Already recent work has shown how the use of EFTs can be adapted to standard nuclear shell model techniques (Stetcu et al. 2007). The impact of extreme scale computing on LQCD and its implications for the nuclear forces extends well beyond these areas, and provides key

inputs into nuclear structure and reactions, and astrophysics. Already, LEFT calculations have probed infinite neutron matter at densities below that of nuclear matter (Borasoy et al. 2007, 2008).

The following sections outline highlights of recent calculations within this subfield of nuclear physics, many of which required several sustained teraflop-years of computation.

Nucleon Axial Charge

The nucleon axial charge, g_A, is a fundamental measure of nucleon structure determined in neutron β-decay. Its value largely dictates the rate of proton-proton fusion, the first step in the thermonuclear reaction chains that power low mass hydrogen burning stars like the sun. Presently, its value is empirically constrained at the 0.2% level. Thus, a confrontation of a precise LQCD calculation of g_A with experiment provides a benchmark for the ability to calculate hadron structure, as well as directly impacting, and more importantly, improving standard solar model input physics. Using chiral perturbation theory, analytic expressions for the mass and volume dependence were used to extrapolate present calculations of g_A to the physical pion mass and to infinite volume, obtaining the axial charge with approximately 10% uncertainties, and in agreement with experiment (Edwards et al. 2006a; Khan et al. 2006).

Isovector Nucleon Form Factors

Nucleons are not point-like structures, but rather have a size of approximately 10^{-15} m (~1 fm). How a nucleon's charge and currents are distributed within its volume are encapsulated within its "electromagnetic form factors." Their experimental measurement remains the subject of intense effort. Recent LQCD calculations with successively lighter quarks dramatically show the emergence of a cloud of pions at the periphery of the nucleon (Edwards et al. 2006b; Göckeler et al. 2007), which is consistent with expectations. In particular, the F_1 form factor monotonically decreases toward the experimental result as the light-quark masses used in the calculation are decreased toward their physical values, and a chiral extrapolation of the results encompasses the experimental value. Similarly, as the light-quark masses decrease toward their physical values, these calculations approach the experimental result for the form factor ratio F_1/F_2, measured by polarization transfer experiments such as those at the Jefferson Laboratory in Virginia.

Quark Contribution to Spin of the Nucleon

Nucleons possess angular momentum called "spin," as do quarks and gluons. Knowing how the spin of the nucleon is partitioned between the spin of the quarks, the spin of the gluons, and their orbital angular momentum is central to understanding the nucleon structure. The contribution of the spin of the up and down quarks to the spin of the nucleon has been a central focus of experimental programs, and the methodology for computing the spin within LQCD is well established. The introduction of Generalized Parton Distributions (GPDs) (Müller et al. 1994; Ji 1997a; Radyushkin 1997) affords a means of calculating the contribution of the quarks orbital angular momentum to nucleon spin (Mathur et al. 2000; Brommel et al. 2007; Hägler et al. 2008) through Ji's sum rule (Ji 1997b). Figure 1 illustrates a recent calculation of the dominant contributions of quark spin and quark orbital angular momentum to the spin of the nucleon (Hägler et al. 2008). While the total orbital angular momentum carried by the quarks within the nucleon is small, the amount carried by the up (u) and down (d) quarks separately is, in fact, substantial.

Low-Lying Baryon Spectrum

Powerful theoretical methods have been developed for the analysis of correlation functions produced in LQCD calculations, and have been used in present efforts to determine the excited state spectrum of QCD. For example, within a theory with only gluons, and without quarks, the spectrum of the low-lying bound states of gluons, known as glueballs, has been calculated (Morningstar and Peardon 1997, 1999); their existence is characteristic of the self-interactions between gluons, resulting directly from the non-Abelian nature of QCD. The existence of three quarks bound together to form baryons is also emblematic of QCD. Recently, a calculation of the baryon resonance spectrum was performed in a theory with gluons and two flavors of quarks, and for the first time a spin-5/2-state was identified in a LQCD calculation (Bulava et al. 2009).

Radiative Transitions in Charmonium

Charm quarks are heavier counterparts of the up and down quarks that comprise everyday matter, with a mass exceeding that of the proton. Because the charm quark and antiquark are heavy, their bound states, collectively known as charmonia, provide a theatre in which to explore the forces between quarks and antiquarks at shorter distances than those characteristic of mesons composed of the lighter up and down quarks. Furthermore, there is presently a wealth of high quality experimental data on the properties of charmonia. The calculation and experimental measurement of electromagnetic transitions between different states of charmonia provides a method of discerning their structure. Recently, the first LQCD calculations of the transition form factors between the lowest-lying charmonium states were performed, and showed good agreement with experimental measurements (Dudek et al. 2006), as illustrated in Figure 2. In a novel extension, the two-photon decays were computed (Dudek and Edwards 2006), and together these calculations lay the groundwork for future computations of the transitions and photo-couplings for mesons composed of light quarks.

Meson-Meson Scattering Lengths

The amalgam of EFT and LQCD has enabled rigorous and precise calculations of the s-wave scattering lengths for $\pi^+\pi^+$ (isospin-two) scattering (Beane et al. 2008a), and of $K\pi$ scattering lengths (Beane et al. 2006b); see Figure 3. Recently, three-meson interaction parameters have been measured from calculations involving multi-pion and multi-kaon systems that provided the first QCD calculation of a hadronic three-body interaction (Beane et al. 2008b; Detmold et al. 2008a, 2008b). These calculations also provide the first QCD calculation of pion- and kaon-condensates, which may play an important role in the evolution of core-collapse supernovae.

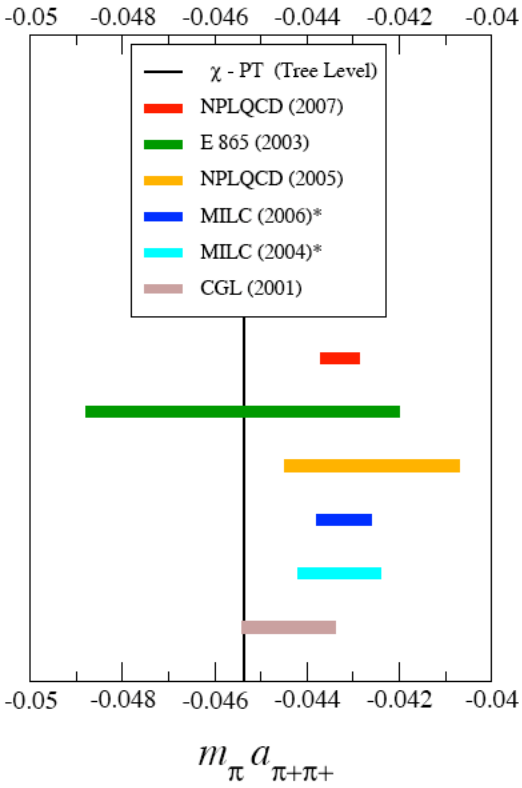

Figure 3. Determinations of the $\pi^+\pi^+$ scattering length. The red bar denotes the high-precision calculation by the Nuclear Physics LQCD (NPLQCD) Collaboration through an amalgam of lattice quantum chromodynamics and chiral-effective field theory. Image courtesy of the NPLQCD Collaboration.

Baryon-Baryon and Meson-Baryon Scattering

LQCD calculations of systems of baryons are less mature because of the poor signal-to-noise ratios encountered in the calculations. Nonetheless, nucleon-nucleon scattering lengths have been calculated, albeit with pion masses too large to provide significant constraints at the physical light-quark masses, demonstrating proof-of-concept (Beane et al. 2006a). The hyperon-nucleon interaction has also been investigated, showing that the scattering phase-shifts for elastic processes—such as $n\Sigma$ scattering—that may be important in the nuclear equation of state in astrophysical settings, can be extracted (Beane et al. 2007).

Development of Effective Field Theories

EFTs provide the bridge between LQCD calculations and other subfields of nuclear physics. Such theories are endowed with a systematic power-counting scheme: at a certain precision of calculation, only a finite number of terms are required in the EFT expansion, and those can be unambiguously identified (Weinberg 1967, 1979). However, there remains low-energy constants that must be determined within these theories; LQCD thus offers a robust method to perform this determination. Furthermore, after these EFTs have been constrained by LQCD calculations, they can be used to investigate processes that may not be calculable with LQCD at that time, such as inelastic processes in nuclear reactions. EFTs can also be used to perform the extrapolations to infinite volume and zero lattice-spacing. EFTs

describing the mesonic sector have been constructed (Rupak and Shoresh 2002; Chen et al. 2007) to address lattice artifacts and discretization effects. The formalism for incorporating EFTs into nuclear structure calculations, such as the NCSM (Navrátil and Ormand 2002) have been (Stetcu et al. 2007) and continue to be investigated.

Lattice-Based Effective Field Theory Calculations

Lattice-based calculations that include nucleons as point-particles and describe their interactions with EFTs have been performed for both pionless (the EFT describing the interactions between nucleons at very low energies) and pionful (chiral) theories. In pionless EFTs, such calculations can probe the universal properties of spin-1/2 particles in a regime where the range of interaction is small compared to the scattering length of the interaction. Such studies (Lee and Schaefer 2005; Lee 2006) have helped the current understanding of dilute neutron matter (at approximately 10% of nuclear matter densities) and have relevance to present day ultra-cold atomic experiments that use magnetically tuned Feshbach resonances and optical lattices. Pionful EFTs have to date probed the spectrum of light nuclei, neutron-nucleus interactions, and the ground state of dilute neutron matter (Borasoy et al. 2007, 2008).

PRIORITY RESEARCH DIRECTIONS

The progression to extreme computing will provide not only increasing precision in present-day calculations, but will open exploration of new opportunities for LQCD to transform scientists' understanding of, and ability to calculate, quantities that are important in nuclear and hadronic physics. The following four PRDs were identified in nuclear physics where extreme computing is essential.

Spectrum of Quantum Chromodynamics

Quarks and gluons combine to form hadrons, known as baryons such as the proton and neutron, and mesons, such as the pion and rho meson. Determining the masses of these hadrons and their excitations from LQCD and comparing with experimental results are vital if scientists are to claim to have a complete description of the theory. Notably, a calculation of the spectrum of meson resonances (those produced from quark and antiquark pairs), and those that have explicit gluon degrees of freedom, may ultimately help understand the mechanism of quark confinement.

How Quantum Chromodynamics Makes a Proton

Nucleons are color-singlet objects composed of quarks and gluons, bound through QCD. The distribution of the quarks and gluons determines how charge, current, mass, and spin are distributed inside a nucleon, and is the subject of intense experimental effort. This research direction aims to answer many outstanding questions related to this issue, and will ultimately build a three-dimensional picture of how quarks combine to make the fundamental building blocks of nuclei.

From Quantum Chromodynamics to Nuclei

Nucleons combine together to form nuclei. Interactions that bind nucleons have their origin in QCD, but their exact lineage presently eludes scientists. This research direction will quantify the connection between QCD and the nuclear forces, determining aspects of the nuclear forces that are poorly constrained empirically. Ultimately, this research will provide greater insight as to how nature forms

nuclei that compose the periodic table of elements, and provide the ability to compute the properties and interactions of nuclear systems that cannot be accessed experimentally.

Fundamental Symmetries

Nature is known to be approximately invariant under certain discrete transformations, such a spatial inversion (P) or motion-reversal (T). It is well known that the combined operation of "charge conjugation," C, and "parity," P, (CP), is violated by the weak interactions. However, there is not enough CP-violation in the weak-interactions alone to explain the matter-antimatter asymmetry observed in the universe. In general, such violations are small, but can have measurable consequences in certain physical processes. This research direction aims to quantify the connection between certain physical observables indicative of discrete symmetry-violation, and the underlying interaction that is responsible.

Spectrum of Quantum Chromodynamics

Computing the bound state spectrum of QCD is vital if scientists are to claim a complete description of the strong interactions, and the confrontation of high-precision calculations of the spectrum with future experimental measurements is a vital test of the theoretical framework. Figure 4 illustrates significant milestones that will impact scientists' understanding of QCD as available computing resources evolve toward the extreme scale. In contrast to electromagnetism, the "field-lines" between a quark and antiquark in QCD do not diffuse over large distances, but rather are confined to compact "flux tubes" connecting the quark and antiquark. Baryons themselves are emblematic of QCD, with the three quarks carrying each of the three color charges of QCD. The outstanding arena that spectroscopy provides for exploring QCD is driving intense experimental studies of the spectrum, primarily excitations of the "glue," or gluonic degrees of freedom, with the GlueX experiment, a flagship component of the 12 GeV upgrade at the Jefferson Laboratory. Extreme computing will provide the *ab initio* theoretical calculations required to capitalize on these experimental investments.

Basic Scientific and Computational Challenges

Unstable Resonances. LQCD provides an *ab initio* method to compute the meson and baryon spectrum by exploiting the finite spatial extent of the gauge-field configurations that are used in such calculations (a "finite volume"). In particular, variation of the lattice spatial-volume provides a mechanism to compute the scattering phase shifts for the resonances and their decay modes. A challenge in the approach to the extreme computing era is the extension of currently established methods for investigating elastic processes to the treatment of inelastic decays, in which there are multiple final states.

Flavor-Singlet Contributions to the Spectrum. Calculations of the spectrum have largely been confined to systems that do not admit the annihilation of an initial-state quark with an initial-state antiquark. The inclusion of such terms will require the calculation of so-called "disconnected contributions" (or "disconnected diagrams") with sufficient precision that the energy spectrum can be resolved. This will necessitate the introduction of improved stochastic estimators or the development of alternative methods.

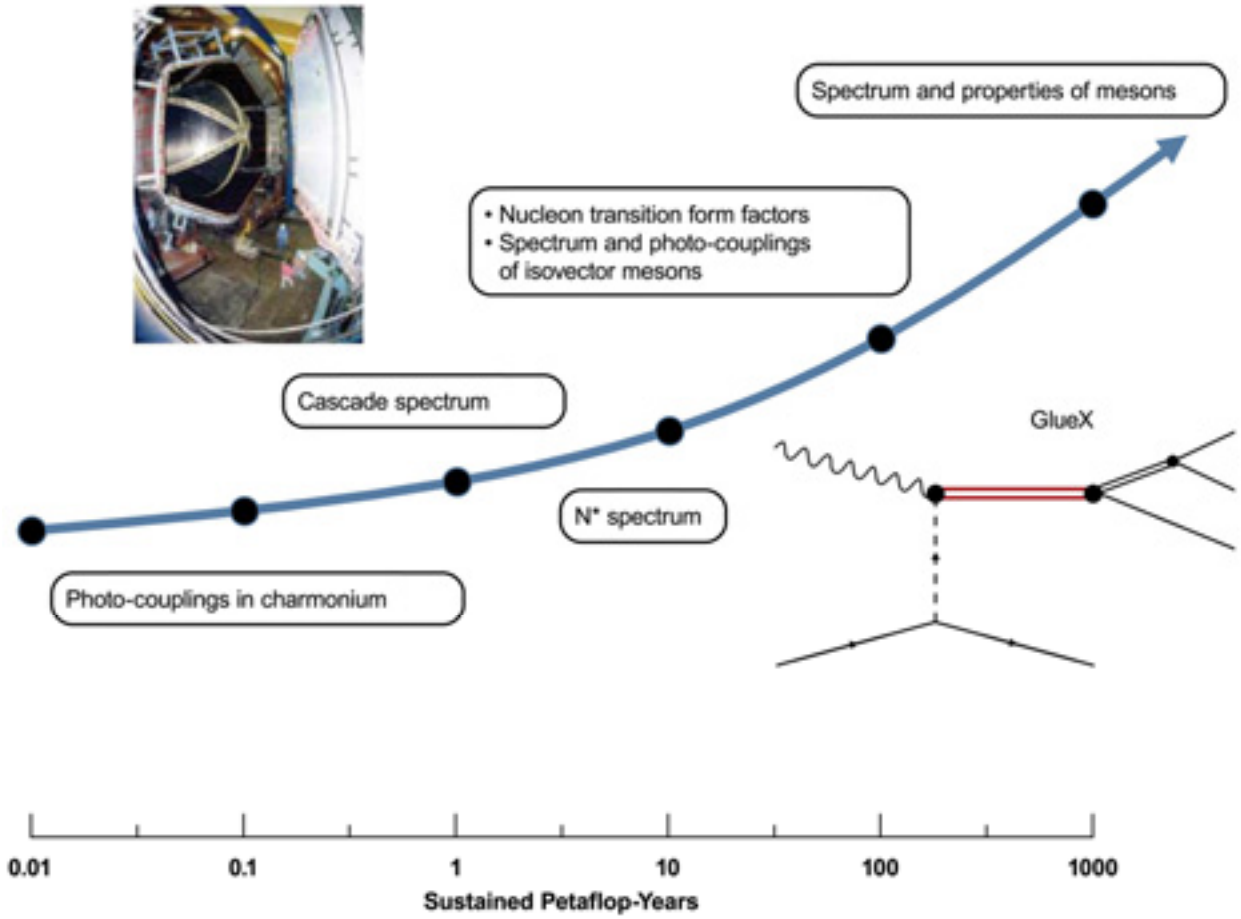

Figure 4. Anticipated highlights for priority research direction "Spectrum of Quantum Chromodynamics." Upper-left image courtesy of the Thomas Jefferson National Accelerator Facility. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Improved Statistical Analysis. As the energy of a state is increased, and as the light-quark masses are decreased, the signal-to-noise ratios of the correlation functions associated with these states generally degrade severely. Current methods provide powerful tools for delineating the different states, but their effective use will require the development of improved statistical tools to fully exploit the investment in leading-edge computing.

Scientific Outcomes and Impacts

The Spectrum and Properties of Meson Resonances. The presently observed spectrum of QCD provides little direct evidence of the presence of gluons. However, QCD presents the possibility of exotic mesonic states of matter in which the gluonic degrees of freedom are explicitly exhibited, and the flux tubes excited. The search for such states will be the subject of intense experimental effort, notably the GlueX experiment at the 12 GeV upgrade at the Jefferson Laboratory (JLab@12GeV). The confrontation of the precise LQCD calculation of the spectrum afforded through extreme computing with the experimentally determined spectrum of meson resonances will provide the culmination of the quest to understand QCD as the theory of strong interactions. The calculation of the spectrum and properties of

exotic resonances will reveal the nature of the gluonic degrees of freedom in the spectrum, and may help elucidate scientists' understanding of the origin of confinement.

The masses and widths of cascade resonances, analogues of the proton and neutron but with two of the *u* and *d* quarks replaced with the heavier, strange quarks, are poorly determined. Even the quantum numbers of many of these states are unknown. Their decay widths are expected to be small and their investigation in LQCD correspondingly less demanding. Computation of the cascade spectrum will require approximately one petaflop-year, and should provide clues as to the role of quark flavor and mass in the spectrum of QCD. Further, these computations are another opportunity for LQCD to provide predictions for future experimental searches.

The spectrum of N* resonances is the subject of intense experimental activity, with its importance encapsulated in the DOE 2009 milestone HP3 (DOE 2008):

> Complete the combined analysis of available data on single π, η, and K photo-production of nucleon resonances and incorporate the analysis of two-pion final states into the coupled-channel analysis of resonances.

The computational methods developed to determine the spectrum of cascades can be extended, but the greater range of decays makes this a more challenging computation, requiring tens of petaflop-years. The baryon spectrum is emblematic of the non-Abelian nature of QCD, and key questions being addressed include the following: what are the roles of the gluons, and more specifically, the role of gluon self-interactions in nucleons? More generally, what are the effective degrees of freedom describing the baryon spectrum?

The experimental measurement of the electromagnetic transitions between low-lying N* resonances are encapsulated in the DOE 2012 milestone HP7 (DOE 2008):

> Measure the electromagnetic excitations of low-lying baryon states (< 2 GeV) and their transition form factors over the range $Q^2 = 0.1 – 7$ GeV^2 and measure the electro- and photo-production of final states with one and two pseudoscalar mesons.

The LQCD calculation of these transitions will require approximately 100 petaflop-years, and provide further clues to the composition of the low-lying baryon spectrum. Furthermore, calculation of the electromagnetic properties with increasing Q^2 (square of the four-momentum transferred to the hadron) enables the perturbative QCD approach to a quark and gluon picture of hadrons to be investigated.

The future GlueX experiment at Jefferson Laboratory's 12 GeV upgrade aims to photo-produce so-called exotic mesons, with the first physics results expected in the middle of the next decade. LQCD has a vital role in both predicting some of the low-lying spectrum, notably for those states with isovector quantum numbers, but also in computing the photo-couplings between these and conventional mesons. These calculations will provide vital input for estimating production rates in the GlueX experiment, and highlights the role of LQCD in guiding experiments.

How Quantum Chromodynamics Makes a Proton

Protons and neutrons, collectively known as nucleons, are the basic building blocks from which all nuclei are constructed, but are themselves formed from the quarks and gluons of QCD. Determining how the quarks and gluons form protons, neutrons, and other hadrons is at the core of frontier nuclear physics experiments at the Brookhaven National Laboratory in New York, the Jefferson Laboratory in Virginia, and international laboratories. Extreme computing is required to perform *ab initio* LQCD calculations of the fundamental properties of nucleons, and provide insight into their structure that is inaccessible to experiment. Together, forefront LQCD calculations and new experimental measurements—such as those exploring transversity and of generalized parton distributions—will enable scientists to build a three-dimensional picture of neutrons and protons in terms of the primordial quarks and gluons of QCD. Finally, scientists will discern how mass, spin, charge, and currents are distributed within a nucleon.

Basic Scientific and Computational Challenges

Calculation of Gluon Contributions to Hadron Structure. Although approximately half the momentum and spin of the nucleon comes from "glue," or gluonic degrees of freedom, calculations of the gluonic contributions within hadrons are far more difficult than those of the corresponding quark contributions. Improved gluonic operators must be developed and the computational infrastructure for much higher statistics calculations will be needed.

Calculation of Flavor-Singlet Gluon and Sea-Quark Contributions. Precision calculations of hadron structure have been largely restricted to *isovector* quantities, such as the difference between the proton and the neutron matrix elements, in which the so-called disconnected contributions cancel and gluons do not contribute. Calculations of proton and neutron properties separately, and more generally the flavor-separated contributions of quarks and gluons to hadron structure, require calculation of disconnected diagrams and their mixing with gluons. Practical calculation of these notoriously difficult quantities will require the development of improved estimators and stochastic noise techniques.

Higher Moments of Structure Functions. Because the ultimate goal is to calculate structure functions and LQCD calculations can only produce moments of these functions, it is desirable to calculate as many moments as possible to optimally reconstruct the relevant physics. Because the lattice has hypercubic symmetry, and not the Lorentz symmetry of the space-time, present techniques only permit calculation of the three lowest moments of the structure functions. Thus, it is necessary to develop new techniques to enable calculation of higher moments.

Form Factors at High Q^2. The ability to determine hadron structure at very short distances, or alternatively at high-momentum transfers, is limited by systematic uncertainties associated with the finite lattice spacing in LQCD calculations, by the degrading signal-to-noise ratios at increasing hadronic energies, and by the decreasing size of the form factors at high momentum.

Figure 5 illustrates the progression toward extreme computing for hadron structure.

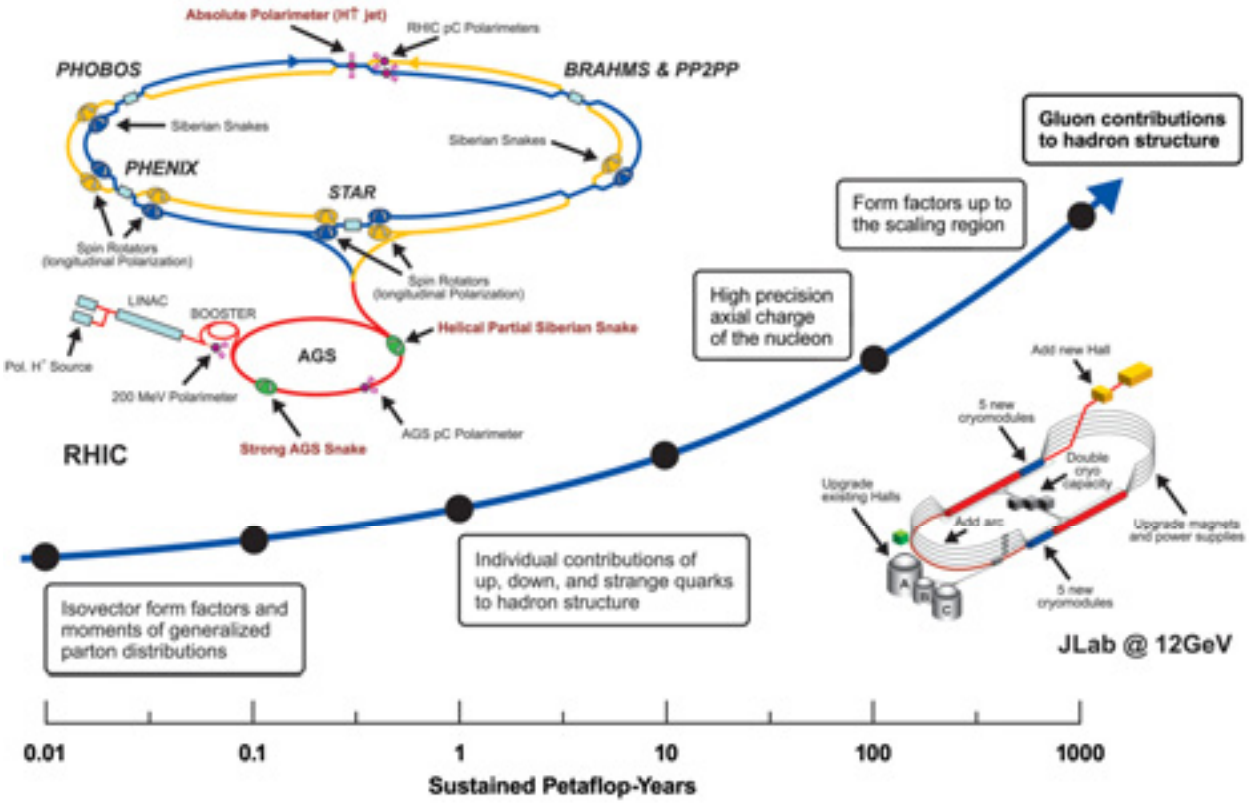

Figure 5. Anticipated highlights for priority research direction "How Quantum Chromodynamics Makes a Proton." Upper-left image courtesy of Brookhaven National Laboratory. Lower-right image courtesy of Thomas Jefferson National Accelerator Facility. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Scientific Outcomes and Impacts

Gluon Contributions to Nucleon Structure. The contribution of gluons to the nucleon mass, and the calculation of the low moments of the spin-averaged and spin-dependent gluon distributions, will address key questions in the 2007 Nuclear Science Long Range Plan (DOE 2007). LQCD calculations are crucial to experimental investigations of the hadron structure of nucleons at the Jefferson Laboratory, Relativistic Heavy Ion Collider-spin and a possible future electron-ion collider. Notably, these calculations, together with experiments, will resolve the origin of spin in the nucleon. These calculations will also delineate between the roles of the spins of the quarks and gluons, and of their orbital angular momentum with a precision that neither experiments nor computation can achieve alone.

The progression toward extreme computing for hadron structure is encapsulated in Figure 5. LQCD will enable precision calculations of key isovector quantities. These include the nucleon axial charge, which impacts the lifetime of the neutron, electromagnetic form factors specifying the spatial distribution of charge and magnetization in the nucleon, moments of quark distributions measured in deep inelastic scattering, and moments of generalized parton distributions, which are a major focus of the experimental program at the Jefferson Laboratory. Calculations requiring computational resources approximately one

petaflop-year at the physical pion mass are required for these observables to be extrapolated to infinite volume and the continuum with an accuracy of a few percent.

Separate calculation of neutron and proton form factors, moments of quark distributions, and of GPDs require the calculation of more computationally demanding disconnected quark contributions, originating from the sea quarks. Using algorithms that have recently been developed, 10-100 petaflop-years will enable calculation of these disconnected diagrams and therefore meet the DOE 2014 milestone HP9 (DOE 2008):

> Perform lattice calculations in full QCD of nucleon form factors, low moments of nucleon structure function and low moments of generalized parton distributions including flavor and spin dependence.

This will include detailed imaging of the two-dimensional transverse spatial structure of the nucleon. An outstanding example of synergy with the experiment is the combination of moments of GPDs calculated with LQCD, and convolutions of GPDs measured at the Jefferson Laboratory and elsewhere, which together will provide a more complete understanding than either effort could separately obtain.

Calculations requiring computational resources of order 100 petaflop-years will increase the precision of the axial charge calculated with LQCD to a level of better than 1%, which will begin to impact the calculation of the proton-proton fusion rate central to solar models.

Extreme computation is required for the calculation of nucleon form factors to sufficiently high-momentum transfer to explore the onset of asymptotic scaling behavior. Such calculations will complement the analogous investigations and calculations of structure in, and transition form factors to, unstable baryons such as the $\Delta(1232)$. The calculation of photon structure functions, hadronic polarizabilities, and the exploration of higher moments of structure functions also requires extreme computing resources.

From Quantum Chromodynamics to Nuclei

In low-energy and low-temperature systems (e.g., conditions as they are on the earth), QCD displays itself through the existence of hadrons (e.g., protons and pions) and their interactions. Exactly how this occurs has been a long-standing question in fundamental physics.

The coupling of EFTs and LQCD (e.g., Beane et al. 2008c) in recent years has allowed for substantial progress in deriving the interactions between hadrons directly from QCD, particularly for systems involving mesons (e.g., pions and kaons). Similar EFTs, when constrained empirically, have been successful in nuclear many-body calculations of light nuclei (e.g., alpha and lithium). LEFTs using nucleon degrees of freedom, as opposed to those of quarks, have made large strides in calculating neutron matter, an infinite medium consisting of neutrons, albeit at small densities compared to normal nuclear-matter densities. These research developments are at a nascent stage because of computational limitations. Extreme computing is required to bring research in these areas to full maturity and become a cohesive program.

How Do Nuclei Come from Quarks and Gluons?

At high energies (>> 1 GeV), or equivalent distances less than 1 fermi (10^{-15} m), quarks are free to move around. However, at lower energies corresponding to distances comparable to 1 fermi or greater, quarks become glued to each other by virtue of their complicated interactions with other quarks and gluons. This phenomenon is known as confinement. At these distances, a quark and antiquark, for example, combine to make a meson, such as a pion or kaon. Exotic forms of mesons come from quark/antiquark pairs with explicit gluonic degrees of freedoms. Three quarks can combine to make baryons. For example, two up and one down quark combine to make a proton, whereas two down and one up quark make a neutron. Certain baryons can then attract each other via the strong nuclear interaction, which roughly can be viewed as a "residual" interaction of quantum chromodynamics. The attraction between nucleons (protons or neutrons) is large enough that groups of nucleons stick together, forming the elements of the periodic table, such as the ^{4}He nucleus shown below.

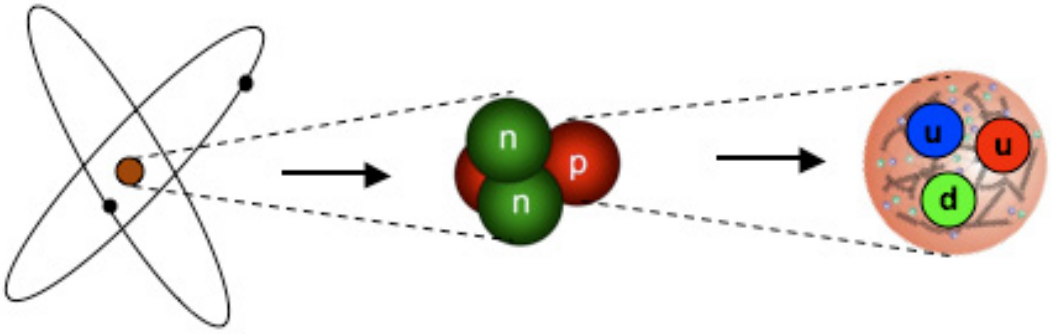

The ^{4}He atom, consists of a nucleus that has two electrons orbiting it (left). The nucleus itself is composed of nucleons–in this case two neutrons and two protons (center). Each individual nucleon, in turn, is composed of three quarks immersed in a mass of virtual gluons and quark/antiquark pairs (right). Image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

For example, with extreme computing resources, LQCD calculations of the interactions within multibaryon systems, like the triton and alpha system, will allow for precise extraction of the three-nucleon interaction, a quantity that is currently poorly constrained empirically. EFTs will be constrained not only empirically, but also by first-principles LQCD calculations, and will be directly fed into many-body nuclear calculations of nuclei, such as those being performed using the NCSM and CC formalisms. These same theories will be used to calculate neutron matter at larger densities, which will help scientists calculate the properties of nuclear matter—for example, in the outer crusts of neutron stars. This PRD will therefore have impact not only at the microscopic level, but to earthly and astrophysical phenomena as well.

What is the Force Between Protons and Neutrons?

The interaction between protons and neutrons can, to a good extent, be described by an "NN potential" (nucleon-nucleon potential) in the same way that the electrostatic interaction between charged particles (such as between two electrons) is described by the Coulomb potential. Nuclear physicists have been able to empirically determine many aspects of NN potentials (see figure below) to high precision and, to date, potentials have been used with great success in nuclear structure calculations of many key features of light nuclei, such as the binding energies of the helium and lithium nucleus. However, as opposed to the Coulomb analogy, there exists the possibility of a "NNN interaction" (three-nucleon interaction), where triplets of neutrons or protons (or combinations) can interact in addition to pair-wise forces. Recent high performance nuclear structure calculations emphasize the need for such a three-body force, but empirical data to constrain this force are lacking. LQCD offers a controlled method for calculating this three-body interaction, given sufficient computational resources. The three-body force, to date, represents the largest uncertainty in nuclear structure and reaction calculations. Precisely determining it from LQCD calculations will greatly improve current understanding of various physical phenomena, from fission processes in carbon-free energy producing nuclear power plants to fusion processes that power the sun.

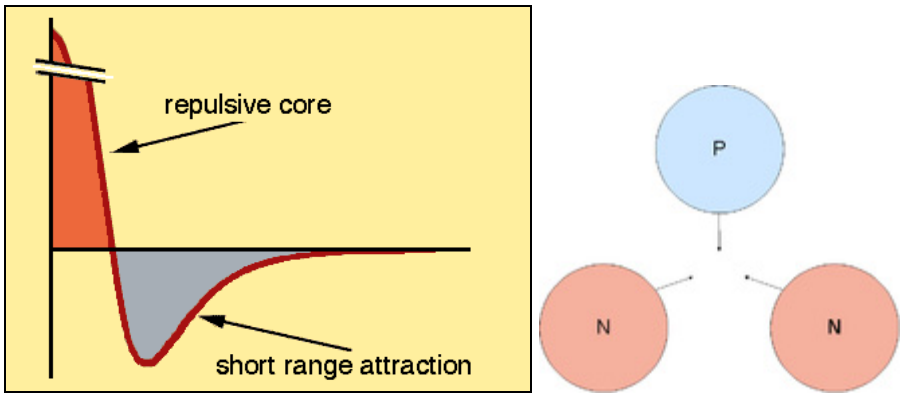

The left figure is a schematic illustration of the potential between a proton and neutron. There is a long range attraction that occurs at large distances and a strong repulsive core at short distances. The right figure schematically shows a three-nucleon interaction. This interaction occurs in addition to any pairwise interactions between the different nucleons. Left image courtesy of Florida State University; right image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Basic Scientific and Computational Challenges

Signal-to-Noise. Current LQCD calculations are inherently stochastic. That is, these calculations use random sampling techniques (Monte Carlo) to perform high-dimensional integrals that are necessary to describe physical phenomena. Calculations of baryon systems, such as the deuteron, suffer from poor signal-to-noise ratios due to the stochastic nature of these calculations. This impedes extractions of multibaryon interaction parameters. As computational resources increase, signal-to-noise issues will diminish slightly, but only with the development of novel algorithms and computational techniques can the signal-to-noise issue be resolved. A similar issue, commonly known as the "fermion sign problem," is seen in LEFT calculations of neutron matter.

Scaling Multi-Baryon Codes for High-Performance Capability. Existing algorithms for performing multibaryon calculations are not suited for extreme scale computing, and will have to be modified to take advantage of extreme computing capability. However, such modifications will not be straightforward, and novel algorithms for multibaryon systems need to be developed to optimize use of large computational resources.

Development of Finite-Volume EFTs. Continued development of finite volume EFTs needs to occur so that LQCD calculations can be matched onto theories used by other areas of nuclear physics, such as nuclear structure and reactions, and nuclear astrophysics. Such theories will need to be "pionful," allowing for the (perturbative) determination of the light-quark mass-dependence of the interactions and scattering parameters. This will provide the most robust extrapolation methods.

Interfacing with Large-Scale Nuclear Structure Calculations. LQCD calculations of few-nucleon interaction parameters will ultimately be fed into nuclear many-body calculations via the use of EFTs, such as those being performed with NCSM and CC theories. EFTs matched to LQCD calculations (typically in a plane-wave basis) will need to be adapted for nuclear structure calculations (which typically use the harmonic oscillator basis). At the two-particle level, this matching is straightforward in that there is an exact analog of Luscher's formalism (Luscher 1986*)* within a harmonic oscillator basis. However, for three- and higher-body nucleon systems, significant research remains to be done. This will entail substantial collaboration with the nuclear structure community—something that is currently just beginning. Both theoretical and numerical methods need to be developed to enhance the overlap between LQCD and the nuclear structure and reactions community. Similar efforts need to be made with the nuclear astrophysics and "hot and dense" QCD scientific communities.

Scientific Outcomes and Impacts

Extreme scale computing will greatly extend the computational prowess of the nuclear physics community. Certain calculations, only aspirations before, will now be accessible and have transformational impact on the broader physics community as a whole. The following are some key scientific outcomes in this PRD that will result from extreme scale computing resources.

Three-Body Interaction Between Baryons. LQCD calculations of three- and four-baryon systems, such as the triton and alpha particle, will allow for the extraction of various three-body interaction parameters that are currently poorly constrained (if at all) empirically. Of particular importance is the three-nucleon interaction, which has implications to the nuclear structure and reactions community. The three-body interactions between nucleons and hyperons will also be accessible, which in turn could have direct implications in astrophysical settings.

Binding Energy of Alpha Particle. For the first time, the four-nucleon system will be calculated directly from QCD. This particle represents the heaviest s-shell nucleus—its inclusion into the suite of LQCD calculations will allow for a comprehensive constraint on the interaction parameters in the EFTs needed for nuclear many-body calculations. Probing this system will also give scientists insight into the four-nucleon interaction—something that presently cannot be done experimentally.

The impact of the above outcomes, and the research leading up to these outcomes, are widespread. For example, insufficient knowledge of the three-nucleon interaction is responsible for the largest systematic uncertainties in nuclear structure and reaction calculations of light nuclei. Without a better knowledge of

this interaction, absolute binding energies and level orderings of excited states of nuclei cannot be calculated with high fidelity. LQCD calculations at the extreme scale will remove this obstacle. Furthermore, research leading to these calculations will help scientists understand the interactions between two- and three-body systems that are not accessible experimentally, but believed to play prominent roles in astrophysical settings (e.g., the interaction between kaons and nucleons). LEFT calculations of neutron matter will be performed at a precision where scientists can quantitatively state the properties of the crust of neutron stars, which are the remnants of Type-II supernovae.

Figure 6 shows anticipated key highlights obtained with high performance computing as extreme computing capability is approached. Baryon-baryon interaction parameters will be computed in the limit of exact isospin symmetry with high precision with sustained petascale resources. With an order of magnitude increase in computational resources, the deuteron axial-charge will be accessible. This is one of the key ingredients constraining certain fusion reactions—within the sun, for example. Finally, at the extreme scale, the three-nucleon interaction can be calculated, as well as the alpha-particle system. Throughout this period, development on EFTs and their overlap with other subfields of nuclear physics will be performed.

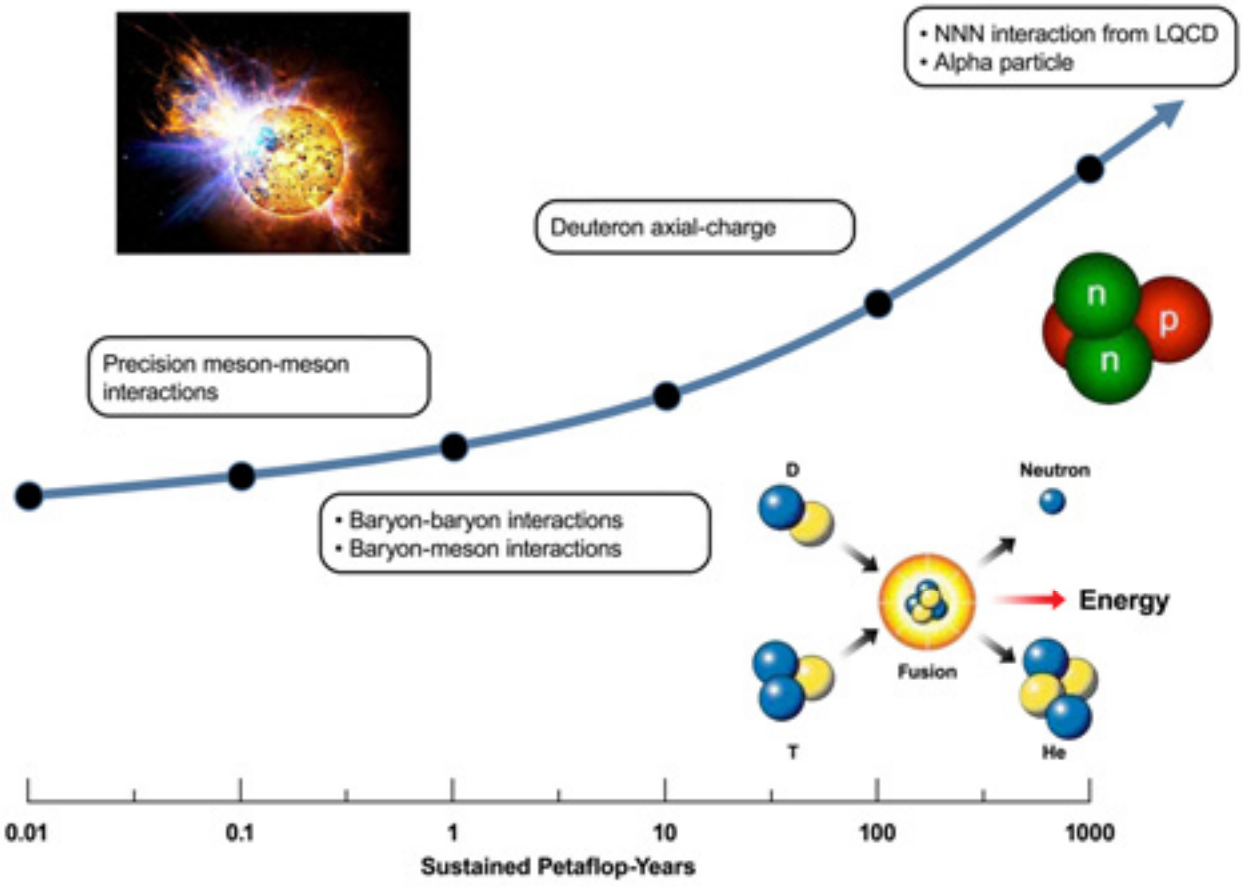

Figure 6. Anticipated highlights for priority research direction "From Quantum Chromodynamics to Nuclei." Upper-left image from NASA. Lower-right image courtesy of the Plasma Physics Laboratory of the Royal Military Academy, EURATOM Association, Belgium. Remaining images courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

The authors of this report emphasize the research impact in this direction will come from calculating observables that are currently inaccessible by experiment, and have great relevance not only to the QCD community, but to the broader nuclear physics community such as those of nuclear structure and nuclear astrophysics.

Fundamental Symmetries

In some instances, nature is very nearly invariant under certain symmetry transformations, such as spatial inversion or motion reversal (also known as time-reversal). However, the consequences of a slight noninvariance under such transformations can have widespread implications. A well known example is CP-violation, where the combined symmetry operation of charge-conjugation, C, and spatial-inversion, P, is known to be slightly violated. Without CP-violation, the present-day matter and antimatter asymmetry of the universe would not exist (the universe contains more matter than antimatter), and from what ensues, humans would not exist.

Research efforts to uncover particles and symmetries beyond those of the standard model of particle physics are multipronged. One of the approaches in this effort is to perform precision measurements of the properties of known particles, such as the magnetic moment of the muon. The E821 experiment at Brookhaven National Laboratory has measured the deviation from the classical value of the muon magnetic moment, g-2, to eight significant digits, and is found to agree with the theoretical calculation within the uncertainties of both the theoretical calculation and the experimental determination. One of the significant uncertainties in the theoretical calculation arises from strong interaction contributions through quantum loops. Exploratory LQCD calculations are underway to understand the methodology that may be employed to directly calculate these loop contributions.

This PRD aims to quantify the connection between certain violations of fundamental symmetries and the resulting observed physics phenomena. Explicit examples include the parity-violating part of the nuclear interaction, and the electric dipole moment (EDM) of the neutron due to time-reversal violation. This will connect to the DOE 2020 milestone F115 (DOE 2008):

> Obtain initial results from an experiment to extend the limit on the electric dipole moment of the neutron by two orders of magnitude.

Extreme computing will allow, for the first time, a quantitative understanding of how these broken symmetries manifest themselves in nuclear physics interactions. Scientists will gain a much deeper understanding of how these symmetries, at the fundamental level and in particular through electroweak interactions and interactions beyond the standard model of particle physics, impact nuclear physics.

What Are Fundamental Discrete Symmetries?

Nature presents various symmetries; i.e., fundamental laws of physics are unaltered by various transformations.

Three of the most classic types of discrete transformations are parity inversion (P) where spatial dimensions are inverted, motion-reversal (T) where time is inverted, and charge conjugation (C) where particles are converted to antiparticles. To the precision with which scientists can perform experiments, it is found that the laws of nature are unchanged under simultaneous transformations of parity, time reversal, and charge conjugation, CPT. However, some of the laws are changed by individual applications of one or two of these transformations. For example, the weak interaction (one of the four known forces of nature) maximally violates both P and C, and neutrinos have only left-handed (their intrinsic spin is aligned with their direction of travel) interactions in the standard model of particle physics. Further, CP-violation in nature is one of the necessary ingredients for having the matter-antimatter asymmetry that is observed in the universe. The figure below shows how the existence of a neutron electric dipole moment (EDM) would indicate the violation of motion-reversal symmetry (T).

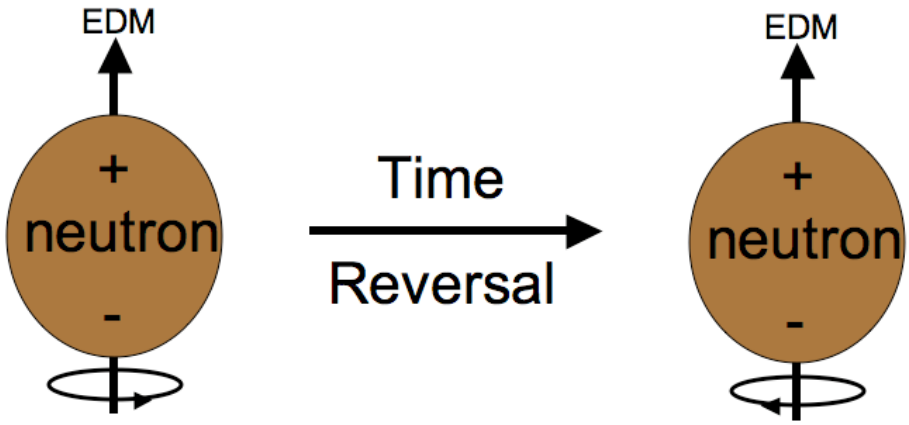

Consider a neutron with a prescribed spin and a (hypothetical) EDM as depicted in the left side of the image. Under a time-reversal transformation, the motion of the neutron is "played backwards," and consequently the neutron spins in the opposite direction, as shown on the right. However, the charge distribution within the neutron that gives rise to the EDM remains unchanged. As the neutron EDM must be proportional to the neutron spin and these two quantities have opposite transformations under motion-reversal, a nonvanishing EDM violates time-reversal symmetry. Quantum chromodynamics allows for time-reversal symmetry violation, and can thus endow a neutron with an EDM. Image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Basic Scientific and Computational Challenges

Four-Point Functions. Calculations of these symmetry-violating observables will generally require a new class of algorithms that enable the calculations of four-point functions. The demands of these calculations will require high-petascale capability, with full maturity coming from extreme scale resources.

Disconnected Diagrams. To date, LQCD calculations have generally involved a certain class of calculations—those of connected diagrams—because of computational limitations. These diagrams represent the propagation of quarks from the initial to the final states. For a complete description of parity violation in nuclear structure, for example, there are short-distance parity-violating few-nucleon forces that contain disconnected diagrams.

Sampling Relevant Topological Sectors. Because symmetry-violations are typically small, the Monte Carlo calculations of these phenomena have signal-to-noise ratios that diminish in time much faster than in most standard LQCD calculations. These "topological fluctuation" issues cannot be remedied by simple reweighting techniques, and will require significant resources to test and develop techniques to avoid this problem.

Memory Requirements. Lattice measurements of all parity-violating effects will have substantial memory requirements because of the large number of distinct light-quark propagators that will be required. Such requirements are currently estimated to be at least two orders of magnitude greater than current available resources.

Scientific Outcomes and Impacts

Parity-Violating Nuclear Interactions. It is generally agreed by scientists that there should be a nucleon-nucleon interaction mediated by one-pion exchange that arises from the weak interaction, and thus a "long-distance" parity-violating contribution to the nuclear force. This parity-violating effect, which is encoded in the weak analogue of the nucleon axial coupling, remains poorly determined despite decades of experimental effort. A LQCD determination of this coupling will have a great deal of impact. In principle, all parity-violating effects in the two-nucleon sector can be calculated with LQCD by extracting parity-violating two-nucleon scattering parameters from the energy levels of two nucleons in a finite-volume lattice, which interact through four-quark operators. Knowledge of the microscopic origins of parity violation in nuclear physics will help correlate and explain the parity-violating signatures observed in nuclear structure. Moreover, knowledge of the parity-violating nuclear interaction calculated with LQCD will provide an explanation of how the weak interaction at the quark level and the strong interaction conspire to generate weak interaction forces among nucleons, and parity violation in nuclear structure.

Neutron EDM Due to the θ-Term and Higher-Dimension Operators. It is possible that QCD contains CP-violating effects that propagate into the hadronic sector via the so-called θ-term (and also through higher-dimension, "irrelevant" operators). One approach to isolating and quantifying these effects is a direct LQCD measurement of the neutron EDM with a nonzero value of θ. While there have been some preliminary studies, this is a computationally challenging endeavor as the CP-violating effect is expected to be small and its signal quickly diminishes in time as topological fluctuations become smaller in the approach to the chiral limit. An understanding of the presence of CP violation due to the θ-term in QCD will sharpen the search for CP-violation whose origin is beyond the standard model, and more generally, constrains models of physics beyond the standard model.

Figure 7 shows the anticipated milestones of this PRD as the extreme scale computing era is approached. Preliminary calculations of the hadronic parity-violating part of the nuclear interaction will be obtained with petaflop-year sustained resources, followed by the first calculations of the neutron EDM with an order of magnitude increase in computer resources. At the extreme scale, the full effects of the nuclear parity-violating component of the nuclear interaction, as well as the full understanding of the neutron EDM due to the θ-term and higher-dimension operators, will come to fruition. These results will coincide with measurements obtained with experiments at the Spallation Neutron Source at Oak Ridge National Laboratory.

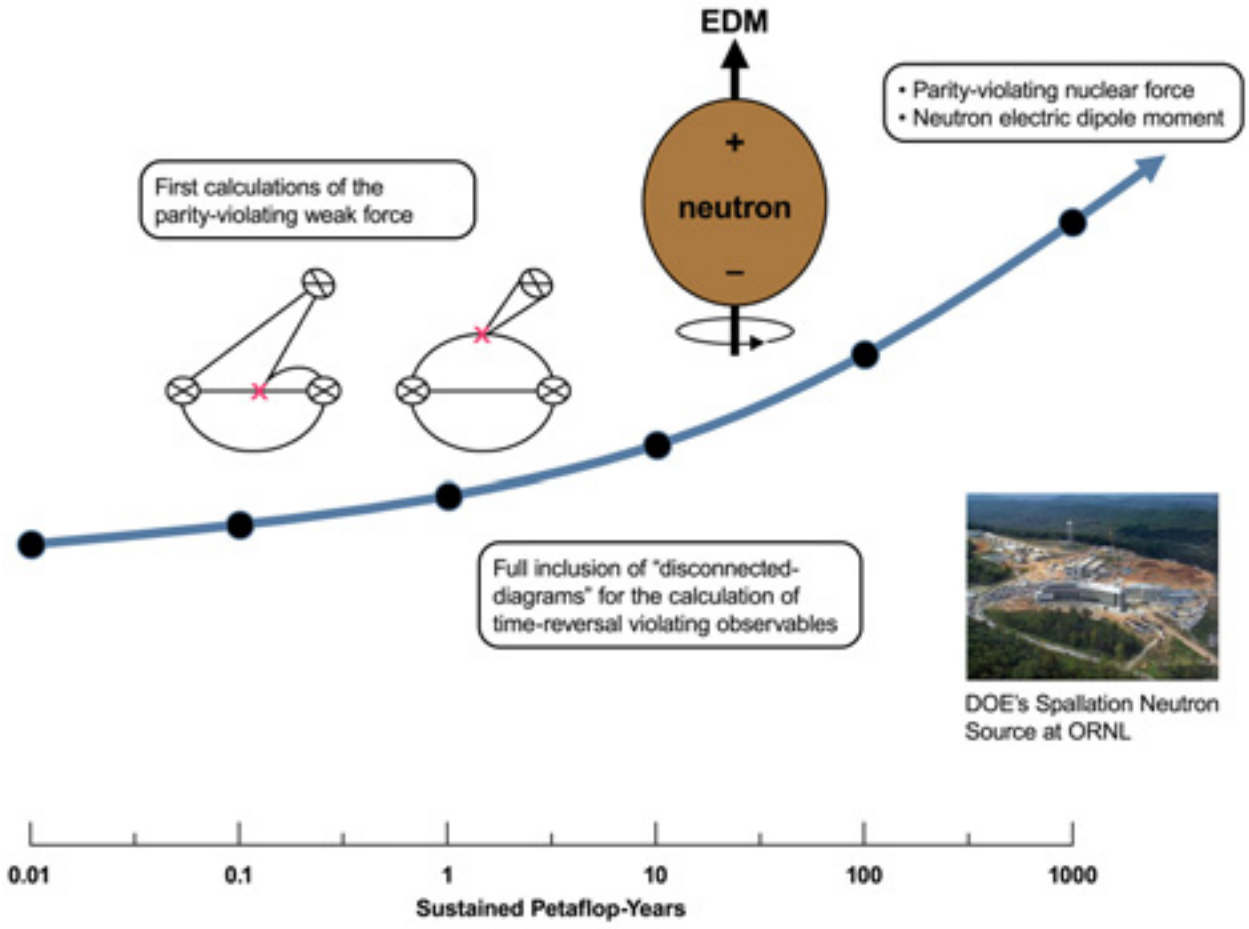

Figure 7. Anticipated highlights for priority research direction "Fundamental Symmetries." Lower-right image courtesy of U.S. Department of Energy, Oak Ridge National Laboratory. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

APPLIED MATHEMATICS AND COMPUTATIONAL ISSUES

At the center of many computations in LQCD is the solution of a large system of linear equations, arising from the lattice-Dirac equation. Typically, this is of the form as shown in Equation (1):

$$Mx = b \tag{1}$$

where M is the Dirac operator, and b represents a large number of right-hand sides. In addition to the high dimension and large number of right-hand sides, the matrix M can become highly ill-conditioned (known as critical slowdown). Currently, scientists rely on deflation techniques using Krylov-based methods, which have been essential for substantially reducing the computational cost of measurements. For future large-scale computations, domain decomposition or multigrid methods offer more advantages because of the superior algorithmic complexity. However, both of these approaches can be difficult to parallelize efficiently beyond more than a few thousand processors and will require substantial research to reach good parallel efficiencies for extreme scale systems.

Research is also needed on developing new approaches for effective preconditioning of the Dirac operator, particularly in preconditioners that will work efficiently on the proposed new machines. It is very clear that porting current algorithms to the new multicore machines will not be the optimal way to exploit all of the available parallelism at the extreme scale level. Other options, such as parallel sparse direct solvers for use in preconditioners, may also need to be explored because these methods can usually achieve good speedups for large numbers of processors. However, memory requirements for the input matrix could become prohibitive and may no longer fit into the memory of any one processing node, necessitating further research in this area.

Because the final products of LQCD calculations are hadronic correlation functions, a reconsideration of the methods of their calculation from quark propagators is required before moving to extreme scale computing. The current model of computing quark propagators and then performing contractions has its limitations when of the order of 10^3 correlation functions need to be constructed, and when those correlation functions involve increasing numbers of terms. Furthermore, correlation functions made from billions of terms will pose a major challenge in computing them to the necessary precision, potentially inhibiting multibaryon-system calculations. This problem is not isolated to LQCD; the need for extended and arbitrary high-precision arithmetic has recently been demonstrated across a wide span of scientific problems (Bailey 2008). These extended precision arithmetic methods and the associated software, Arbitrary Precision (ARPREC), have already been used in certain contexts in high-energy physics and could find wider use as the calculations and data become even larger. A representative example is given by scientists in the Nuclear Physics LQCD (NPLQCD) Collaboration who have published some of their findings in Beane et al. (2008b) and Detmold et al. (2008a, 2008b). In these reports, the authors discuss the failure of using double-precision calculations when computing the quark propagator contractions required for the correlation functions of systems comprised of more than nine mesons. However, by using 64-digit precision arithmetic, Beane et al. (2008b) and Detmold et al. (2008a, 2008b) were able to compute the correct results. These numerical issues will particularly impact calculations of future multinucleon systems, thereby requiring more detailed studies of the effects of extended precision arithmetic on LQCD computations.

Another research area is the use of the Monte Carlo as a solution method for computing high-dimensional integrals. Monte Carlo methods in physics are synonymous with Markov Chain methods for sound reasons: the systems that must be studied have far too many states to list, and the fraction of significant states is too small for chainless sampling methods. Markov Chain methods offer intelligent searches, in which states are sampled one by one, by strategies that seek out high-likelihood states. The price is that consecutive states are not independent, so that multiple steps are needed to obtain one new independent sample. Some recent work by Chorin (2008) on a novel method for Monte Carlo sampling could be applicable to nuclear physics problems; for example, hadron formation and quark confinement. At the center of the construction is the fast evaluation of marginals that can be viewed as a renormalization in the sense of the works of Kadanoff and of Wilson. Chorin (2008) has shown that this method yields good results in several dimensions, with errors small enough to be compensated for by a differential weighting of the samples. Further research will be required to evaluate its full potential for the high-dimensional problems in QCD.

As a concrete example of present-day computational requirements for LQCD calculations, the production of one ensemble of 1000 anisotropic gauge-field configurations (from 11,000 trajectories) with dynamical degenerate up and down quarks and a dynamical strange quark with masses chosen to produce a pion with a mass of approximately 270 MeV, all described by the improved-clover action, and with 28 sites in each

space-direction and 256 sites in the time-direction, requires approximately 10 teraflop-years. The lattice spacing of these configurations is approximately 0.12 fm in the space directions, 0.034 fm in the time-direction, and the extent of each spatial direction is approximately 3.4 fm. Approximately 18 teraflop-years are required to generate 500,000 light-quark propagators that are used to form the correlation functions, and 4 teraflop-years are required to contract the light-quark propagators together to form simple mesonic, one baryon, two-baryon and a small number of three-baryon correlation functions. A total of approximately 31 sustained teraflop-years are required to perform this calculation, which is in a relatively small volume, at unphysically large light-quark masses, and at relatively coarse lattice spacing. Reducing the lattice spacing in the spatial direction at fixed volume by a factor of 2 increases the resource requirements for the production of the gauge-field configurations by a factor of 64; similarly, increasing the size of the spatial extent (at fixed lattice spacing) by a factor of 2 increases the resource requirements by a factor of 12. Reducing the mass of the pion from the unphysical value of 270 MeV down to its observed value of 139.57 MeV (the charged pion mass) requires another factor of approximately 4 increase in computational resources. If instead of the improved-clover action (which does not respect the approximate global chiral symmetry of continuum QCD due to the non-zero lattice spacing), the domain-wall action or the overlap action (which do respect the approximate chiral symmetry) is used to describe the light quarks, the computational resources required to generate the gauge-field configurations and the light-quark propagators in this calculation are increased by more than ten-fold. For the calculation outlined above, the Dirac operator is a (approximately) 68×10^6 times 68×10^6 sparse matrix.

Post-Processing Facilities. An emerging feature of LQCD calculations is the increasing proportion of computational cost devoted to physics measurements rather than the generation of gauge configurations performed at capability computing facilities. While a decade ago less than 20% of the computing capacity was needed for physics measurements on the products of the highest end calculations, today 50% of the total capacity is used for such measurements and the trend continues upwards. The analysis computing facilities need to have performance characteristics at about 1% to 10% of the scale of those done on the capability machines, with an aggregate capacity (for LQCD) of at least the same magnitude as for configuration generation. Thus, in addition to computational resources at the extreme scale, scientists need to ensure a commensurate increase in computational capacity for analysis (Branscomb's pyramid, National Science Foundation Blue Ribbon Panel on High Performance Computing [Branscomb et al. 1993]), as well as increases in network and storage capacity. Large-scale computing systems by current standards will be needed for post-processing of these data produced by the extreme scale machines, and for testing and development of new software.

Robust Temporary and Long-Term Storage. The storage of gauge-field configurations alone at the extreme scale will create O(1)-O(100) petabytes of data per ensemble. The storage requirements of the analysis part of the calculation (second tier of the pyramid) are likely to be an order of magnitude larger than those for the gauge configurations themselves.

The challenges ahead provide good opportunities for interdisciplinary research with applied mathematicians and computer scientists to help establish the groundwork needed for solving these problems. Such interactions have already started and have had an impact on both machine design and algorithm development. Access to large computing facilities now and investment in interdisciplinary collaborations is essential to further progress in extreme scale computing for nuclear physics.

ENHANCED SYNERGISM WITH OTHER SUBFIELDS OF NUCLEAR PHYSICS AND THE BROADER PHYSICS COMMUNITY

The increased computational resources of the past decade have brought various subfields of nuclear physics to a point where interdisciplinary research is becoming a reality. No longer will different subfields of nuclear physics exist independently of one another, as the level of precision of calculations has reached a stage where continued improvement involves direct collaboration with other areas of physics. Nowhere is this more clearly seen than in the area of nuclear forces and cold QCD. For example, in nuclear structure and reactions, the need for a consistent determination of the three-nucleon force can only be satisfied from LQCD calculations of this force. Aspects of the nuclear equation of state used in astrophysics simulations can only be determined through LQCD investigations of exotic nuclear systems.

Many of the calculations within the purview of nuclear forces and cold QCD will also have direct impacts on other areas of physics, such as the standard model of particle physics; e.g., EDM calculations, or fusion through the precise calculation of g_A and the deuteron axial charge, as well as atomic experiments using calculations of dilute fermions at the unitary limit. Indeed, LQCD is vital to the high-energy physics community where the emphasis is on calculating and removing the strong-interaction effects in hadrons to reveal other fundamental interactions, such as the electroweak interactions and possible interactions beyond the standard model, acting on the quarks and gluons. There is a strong symbiosis between the LQCD efforts in the two communities. This includes not only the exploitation of new algorithms and theoretical advances, but also the development of software and computational infrastructure. Sharing the gauge-field configurations has been vital to fully exploiting the available computational resources. This symbiosis will only strengthen in the extreme computing era.

Recent evidence therefore supports the notion that increased computational resources builds connections between subfields of physics. Figure 8 shows some of these connections. Such connections will be so strong that, as the extreme computing era is approached, it is conceivable that subfields of nuclear physics will merge together—a likely scenario for nuclear structure and nuclear forces via LQCD. The importance of extreme computing on each individual subfield of nuclear physics is clear. Just as important will be the holistic impact of extreme computing both on the entire nuclear physics community, and the broader physics community, made possible through the strengthening of connections between these currently diverse fields.

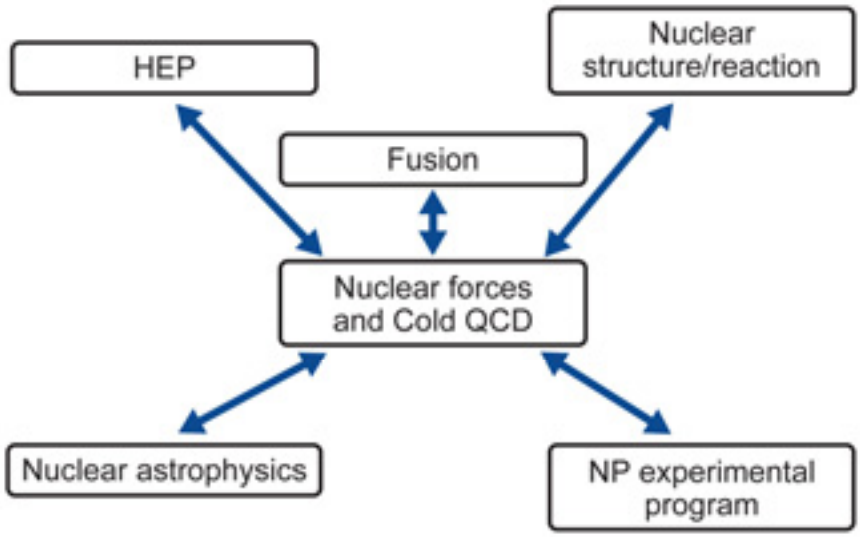

HEP = High-energy physics
NP = Nuclear physics

Figure 8. Extreme computing will drive stronger connections between various subfields of nuclear physics as well as the broader physics community. This figure shows linkages that are expected to become significantly stronger in the era of extreme computing. Image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

CONCLUSIONS

Extreme computing resources are required to reliably connect nuclear physics with QCD, the underlying quantum field theory describing the strong interactions. With extreme computing resources, the properties of and the interactions between protons, neutrons, light-nuclei, and other strongly interacting particles will be reliably calculated directly with QCD, even in environments that cannot and will not be created in any controlled terrestrial experiment.

Future investment in extreme scale computing resources at both the highest tier and at the second tier will provide revolutionary capabilities for performing research in this area of nuclear physics. As detailed in previous sections, the anticipated physics outcomes will be transformational because they impact not only their immediate subfield of nuclear physics, but the broader physics community as a whole.

Feedback from members of the Nuclear Forces and Cold Quantum Chromodynamics subpanel that collaborated during this workshop has been overwhelmingly positive.[1] Science seems to be the major driver for investments in extreme computing, which reaffirms DOE's commitment to providing first-class facilities for performing basic science research.

The progression to extreme computing is required for LQCD to have a major impact on nuclear physics. The amalgam of theoretical, computational, and experimental nuclear physics will transform scientists' understanding of the spectrum, structure, and interaction of hadrons and nuclei, and finally provide a

[1]"Forefront Questions in Nuclear Science and the Role of High Performance Computing," January 26-28, 2009, Gaithersburg, Maryland. Additional information available at http://extremecomputing.labworks.org/nuclearphysics/index.stm.

description of nuclear forces in terms of the fundamental quarks and gluons of QCD. The path to extreme computing will allow scientists to understand how the properties of hadrons and nuclei depend on the fundamental (input) parameters of QCD. The confrontation of LQCD calculations of the spectrum and properties of resonances with experimental investigations such as those being performed (or that will be performed) at national laboratories and elsewhere will enable the low-energy degrees of freedom describing the spectrum to be discerned, and the role of gluons in spectroscopy exhibited. Current experimental investigations of hadron structure at various institutions, combined with LQCD calculations and a rigorous phenomenology program, will allow a three-dimensional picture of the distribution of spin, charge, and matter in the nucleon to be drawn. Furthermore, many longstanding questions within nuclear physics—such as the nature of the three-nucleon or spin-orbit interaction—will be answered, and systems not accessible by experiment can be ascertained via LQCD.

Ultimately, extreme computing will endow LQCD with the ability to provide predictive capability in the realm of low-energy nuclear physics that will have direct implications to nuclear structure, nuclear reactions, and nuclear astrophysics, as well as the broader nuclear physics community.

Table 1 provides an outline of the milestones for the work described in this section. Provided that the computational resources become available for research in cold QCD and nuclear forces at the anticipated scales, the forefront research that will be conducted are provided as milestones.

Table 1. Milestones for Cold Quantum Chromodynamics and Nuclear Forces

Scale	Milestone
>1 Petaflop-year	• LQCD calculations* — Photo-couplings in charmonium — Spectrum of cascade — Isovector form factors and moments of the generalized parton distributions — Precision calculation of meson-meson interactions — First calculations of hadronic parity-violating interactions
>20 Petaflop-years	• LQCD calculations* — Spectrum of excited nucleons — Contributions of up, down and strange quarks to hadron structure — Precision calculation of baryon-baryon and meson-baryon interactions
>100 Petaflop-years	• LQCD calculations* — Nucleon transition form factors — Spectrum and photo-couplings of the isovector mesons — High-precision calculation of the axial charge of the nucleon — Nucleon form factors into the scaling regime — The axial charge of the deuteron, and electroweak interactions — Disconnected diagrams in the calculation of T-odd observables
>1 Exaflop-year	• LQCD calculations — Spectrum and properties of the mesons — Gluon contributions to hadron structure — Three-nucleon interactions, including nnn — Spectrum of the alpha particle — Parity-violating nuclear force — Neutron electric dipole moment — Inclusion of electromagnetism — Physical masses for up and down quarks
*In the isospin limit of QCD without electromagnetism.	

NUCLEAR STRUCTURE AND NUCLEAR REACTIONS

Co-Leads: James P. Vary, Iowa State University
Steven C. Pieper, Argonne National Laboratory

INTRODUCTION AND CURRENT STATUS

Atomic nuclei are the essence of the visible universe. Formed in the big bang or in cataclysmic astrophysical explosions, atomic nuclei are a crucial and intriguing part of the world. The basic features of atomic nuclei were understood in terms of the nuclear shell model in the 1963 Nobel Prize winning research of Eugene Paul Wigner, Maria Goeppert-Mayer, and J. Hans de Jensen. Since then, extensive experimental programs have yielded a detailed knowledge of the nucleon-nucleon interaction. This crucial experimental information will be augmented through studies of quantum chromodynamics (QCD) (see section on Cold QCD and Nuclear Forces). More refined descriptions of nuclei and greater predictive power require understanding nuclear structure and reactions in terms of the underlying interactions. Accurate solutions of these strongly interacting quantum many-body problems will yield new insight into the structure of nuclei and the ability to calculate processes that are difficult or impossible to measure experimentally.

The structure and dynamics of atomic nuclei represent very challenging problems because, unlike the theory of electrons in atoms and molecules, naïve mean-field theories based upon the underlying interactions do not provide even a qualitative description of their structure. The nuclear interaction is strongly spin- and isospin-dependent, and finely tuned to provide a weak nuclear binding of approximately 8 MeV per nucleon, much smaller than a typical scale in QCD.

Large-scale computations have already enabled significant progress in understanding nuclei from the underlying nuclear interactions. The structure and energy levels of the lightest nuclei (consisting of up to 12 nucleons) are well reproduced with realistic two-nucleon interactions plus modest three-nucleon interactions. The same interactions provide good descriptions of electroweak form factors, transitions and response, as well as the nucleon-nucleon correlations as revealed in experiments at Thomas Jefferson National Accelerator Facility (Jefferson Laboratory) and elsewhere. The first low-energy microscopic calculations give confidence that a consistent picture of structure and dynamics is emerging. However, as nuclei with almost 300 nucleons are experimentally observed, studies of systems containing 12 nucleons are only the beginning.

Progress in large-scale nuclear calculations is only now extending to larger nuclei, where much of the promise of extreme scale computing lies. The theory of the atomic nucleus focuses on predicting and explaining rich classes of phenomena that occur in nuclei and nucleonic matter. Atomic nuclei exhibit many intriguing properties; pairing energies, for example, are a significant fraction of the Fermi energy. In heavy neutron-rich nuclei, pairing energies can be comparable to shell closure effects, a regime quite different from the most stable nuclei. The binding energies and electroweak transition rates of these nuclei are crucial to understanding the production of the heaviest elements. Understanding nuclei is also critical in exploring fundamental physics, including for example the absolute mass scale of neutrinos as to be determined in neutrinoless double beta decay.

Given its intriguing nature and importance in the birth of the universe, in astrophysical settings, in energy generation, and in industrial and medical applications, it is fundamentally important and of great practical significance that scientists have a detailed understanding of this complex quantum many-body system. The theoretical goal of increased predictive power for describing nuclear processes that occur in nature or in nuclear reactors, but cannot be measured in the laboratory with sufficient precision, drives scientists in this field to achieve detailed simulations using extreme scale computers and cutting-edge algorithms.

Three major challenges are addressed in this section: 1) a strong inter-nucleon interaction based upon—but not derivable with the necessary accuracy from—QCD; 2) the quantum many-body problem; and 3) phenomena on scales stretching over orders of magnitude in length or energy. Together, these challenges create a computationally difficult problem, and research teams have recently developed a suite of new and sophisticated computational tools capable of addressing previously unsolved problems. A driving force for these teams has been the substantial increase in computational power in the past decade. Consequently, increased computational power has, for the first time, made theories that were known for decades applicable for detailed nuclear physics investigations. The advent of extreme scale computing will enable this progress to accelerate, as shown in this report.

Scientists are on the verge of precision predictions for low-energy nuclear processes of both fundamental and practical significance such as neutron-nucleus and proton-nucleus reactions, electromagnetic transition rates, fusion of light nuclei, fission of heavy nuclei, and structure of short-lived isotopes. Scientists also aim to develop and validate novel approaches to the structure and reactions of heavier nuclei by predicting a universal nuclear energy density functional and solving density-functional theory for these systems. Building on recent successes with leadership-class computers, extreme scale computing will be critical to the advances projected over the next decade.

The following subsections highlight four important questions in nuclear structure and reactions that extreme scale computing, coupled with continued algorithmic improvements, will answer. These questions cover fundamental nuclear phenomena and span the entire chart of nuclei from the light but fundamental nucleus ^{12}C, through intermediate-mass nuclei such as ^{48}Ca, ^{76}Ge, and ^{130}Te, to the fission of the heaviest nuclei, and, finally, beyond the known nuclei at the limit of nuclear existence for the heaviest elements, and at the neutron drip line and neutron matter. The range of these questions is truly comprehensive in scope and an array of advanced techniques—which are well matched to the power of emerging extreme scale facilities—will be required. The principal benefit of this research direction will be the development of a truly predictive capability for the entire periodic table.

PRIORITY RESEARCH DIRECTIONS

Ab Initio Calculations of Light Nuclei and Their Reactions

Basic Scientific and Computational Challenges

A realistic *ab initio* approach to light nuclei with predictive power must have the capability to describe bound states, unbound resonances, and scattering states within a unified framework. Over the past decade, significant progress has been made in understanding the bound states of light nuclei starting from realistic nucleon-nucleon (NN) plus three-nucleon (NNN) interactions (Pieper and Wiringa 2001; Navrátil et al. 2000, 2007; Hagen et al. 2007b). The solution of the nuclear many-body problem is even more complex when scattering or nuclear reactions are considered. For few-nucleon systems (A=2-4),

accurate methods solve the bound state and the scattering problems. However, *ab initio* calculations for scattering processes involving more than four nucleons are still the exception (Nollett et al. 2007; Quaglioni and Navrátil 2008; Hagen et al. 2007a) rather than the rule. The development of an *ab initio* theory of low-energy nuclear reactions on light nuclei is key to further refining scientists' understanding of the fundamental interactions between the constituent nucleons. At the same time, such a theory is required to make accurate predictions of nuclear astrophysics' crucial reaction rates that are difficult or even impossible to measure experimentally. This section highlights a key direction that *ab initio* methods will pursue with exascale resources.

Reactions That Made Us: Triple-Alpha Process and $^{12}C(\alpha,\gamma)^{16}O$

Extreme scale computing will enable the first precise calculation of $2\alpha(\alpha, \gamma)^{12}C$ and $^{12}C(\alpha,\gamma)^{16}O$ rates for stellar burning (see Figure 9); these reactions are critical building blocks to life, and their importance is highlighted by the fact that a quantitative understanding of them is a 2010 U.S. Department of Energy milestone (DOE 2007). The thermonuclear reaction rates of alpha-capture on ^{8}Be (2α-resonance) and ^{12}C during the stellar helium burning (see Figure 9 for a schematic depiction) determine the carbon-to-oxygen ratio with broad consequences for the production of all elements made in subsequent burning stages of carbon, neon, oxygen, and silicon. These rates also determine the sizes of the iron cores formed in Type II supernovae (Brown et al. 2001; Woosley et al. 2002), and thus the ultimate fate of the collapsed remnant into either a neutron star or a black hole. Therefore, the ability to accurately model stellar evolution and nucleosynthesis is highly dependent on a detailed knowledge of these two reactions, which is currently far from sufficient.

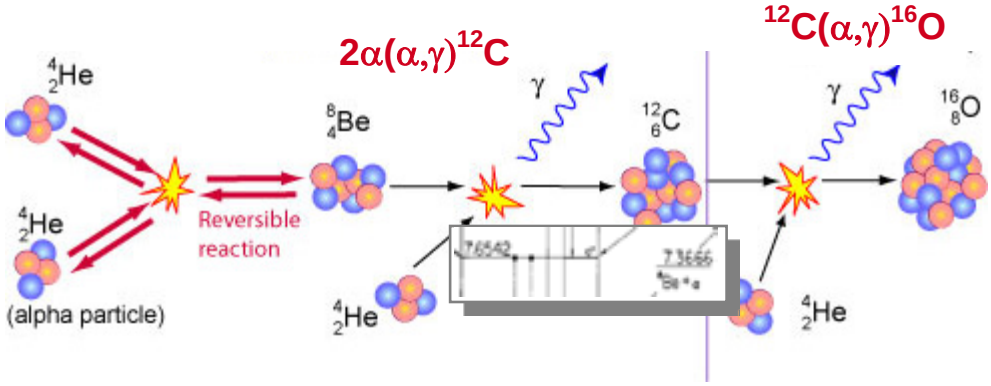

Figure 9. A schematic view of the ^{12}C and ^{16}O production by alpha burning. The $^{8}Be+\alpha$ reaction proceeds dominantly through the 7.65 MeV triple-alpha resonance in ^{12}C (the Hoyle state). Both sub- and above-threshold ^{16}O resonances play a role in the $^{12}C(\alpha,\gamma)^{16}O$ capture reaction. Image courtesy of Commonwealth Scientific and Industrial Research Organisation (CSIRO).

Experimental measurement of these reaction rates at energies relevant for astrophysics (at approximately 300 keV in the center of mass) is impossible with existing techniques because of their extremely small cross-sections. Because of the influence of alpha-cluster resonances in ^{12}C and ^{16}O, theoretical extrapolations of measurements performed at higher energies to the relevant low-energy region have large uncertainties (for recent measurements, see Assunção et al. 2006). Presently, all realistic theoretical models fail to describe the alpha-cluster states, and no fundamental theory of these reactions exists. Yet,

a fundamental theory is needed to determine the rate of the $^{12}C(\alpha,\gamma)^{16}O$ reaction to at least 10% accuracy to fix the subsequent burning stages.

These calculations can be performed by using several independent *ab initio* methods, which will permit results verification and allow for systematic uncertainties to be determined. The methods are as follows: 1) the Green's Function Monte Carlo (GFMC) approach generalized for scattering; 2) the *ab initio* no-core shell model (NCSM) extended by the resonating group method (RGM); and 3) the coupled cluster method. The calculations can proceed in several phases with increasing complexity, and a general picture of the computational requirements for these calculations is shown in Figure 10.

The first phase focuses on the Hoyle state in ^{12}C. This is an alpha-cluster-dominated, 0^+ excited state lying just above the $^{8}Be+\alpha$ threshold and is responsible for the dramatic speedup in the ^{12}C production rate. The calculation of this state will be the first exact description of an alpha-cluster state. It can be achieved with 10% accuracy of the excitation energy within 3 years using the current petaflop machines, and with 5% accuracy in 10 years using improved Hamiltonians. Calculations of alpha-capture on ^{8}Be will be performed within the next 5 years. Calculations for ^{16}O, and in particular of the alpha-cluster resonances that impact the $^{12}C(\alpha,\gamma)^{16}O$ reaction, will follow. Finally, the $^{12}C(\alpha,\gamma)^{16}O$ calculations will be completed within a 10-year time frame.

Scientists can reliably estimate the increase in computer resources needed to address the 16O nucleus with GFMC. Presently, the GFMC calculation of the 12C ground state requires approximately 400 peta operations. The Hoyle state will require tens of calculations of the same size. The number of operations will increase by a factor of approximately 1200 for 16O, with the growth provided by the available computing resources increasing from the petascale to the extreme scale.

Currently, it is becoming feasible to calculate, within the NCSM/RGM framework, low-energy nucleon-^{12}C or nucleon-^{16}O scattering with soft NN forces using approximately 1000 cores on present-day machines. The computational demand increases dramatically (a factor of approximately 10^6) with increasing size of the projectile (from a single nucleon to an alpha particle) and by including the NNN interaction. Therefore, this is clearly a problem requiring the extreme scale computation level.

The ground state of ^{16}O can presently be computed within the coupled-cluster 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 than this because of the proximity of the scattering continuum; it will be based 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 not be easy to use an extreme scale computer for these calculations. The present ability in GFMC was obtained by splitting the work on one Monte Carlo configuration among tens of cores (previously just one core was used). For ^{16}O, 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.

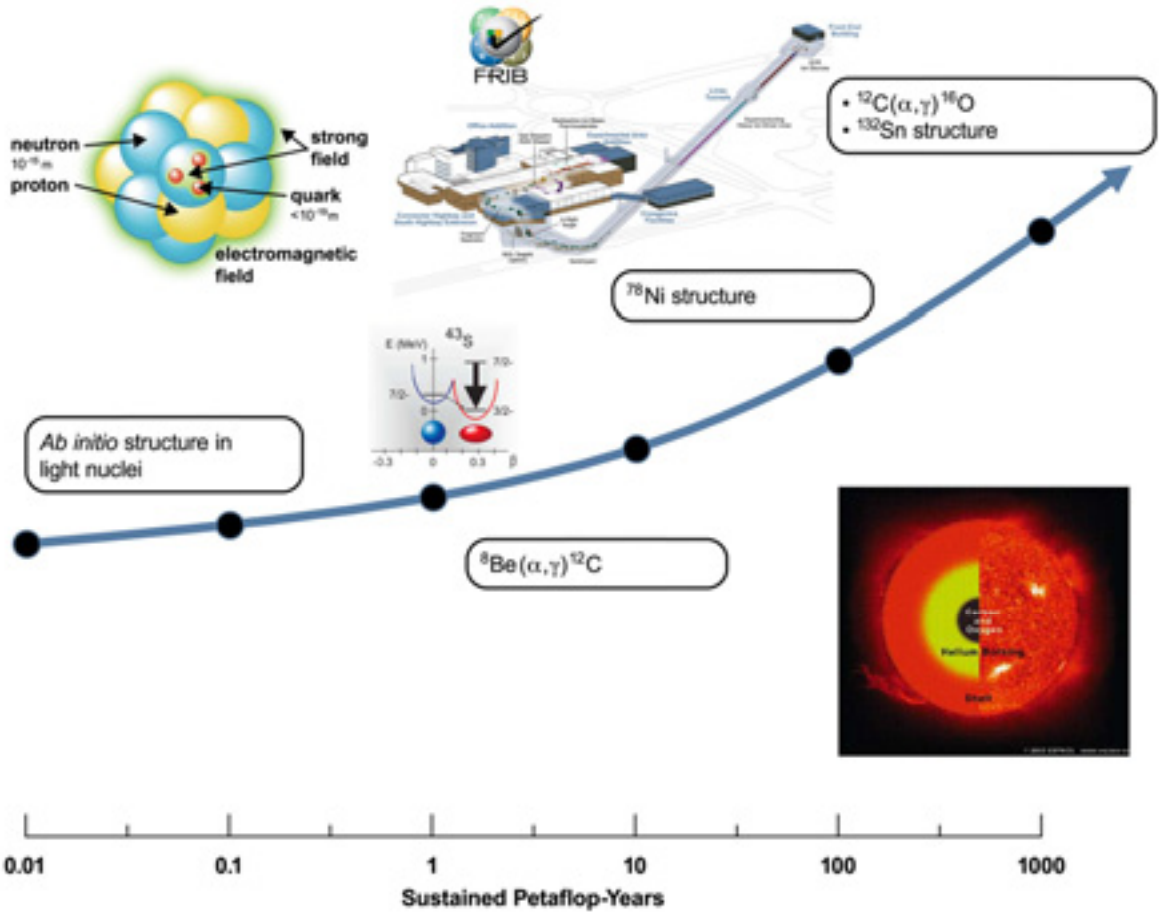

Figure 10. Anticipated highlights for the priority research direction "Reactions That Made Us." Top-middle image courtesy of Michigan State University. Remainder of image courtesy of James P. Vary (Iowa State University).

In the NCSM/RGM, the matrix elements for hundreds of density operators must be calculated. These calculations are both central processing unit and memory intensive. Calculations are presently completed using message-passing interface (MPI) with distribution of the memory allocation. For example, in the calculation of matrix elements of the density operators, the cores are divided into groups, each of which is responsible for computing matrix elements of a subset of operators. This type of parallelization will need to be optimized and propagated to a finer level of distribution among clusters of computing cores in extreme scale machines as the complexity of the task grows rapidly with the mass of the target nucleus, the mass of the projectile (alpha particle), and presence of the NNN force.

Scientific Outcomes and Impacts

The primary outcome of this effort will be a comprehensive understanding of the mechanism behind these two key reactions, and the ability to model the chemical evolution of the universe. Success will permit an accurate determination of the reaction rates at low energies relevant to stellar burning, which are currently limited by large experimental uncertainties. In particular, the uncertainty in the $^{12}C(\alpha,\gamma)^{16}O$ reaction rate is currently about 40% (Angulo et al. 1999). By achieving this research goal, scientists will enhance the predictive power of stellar modeling. At the same time, scientists will develop *ab initio* tools to describe the structure of weakly bound nuclei that will be studied at the Facility for Rare Isotope Beams (FRIB) at

Michigan State University. Verification of model predictions by experiments at FRIB will provide necessary checks on the theoretical approaches and the underlying two- and three-body forces used. One computational outcome will be the development of a library for distributing shared-memory work to subsets of nodes within a massively parallel machine.

A successful completion of this program will provide essential input for modeling of stellar evolution and element production. It will provide a firm basis for extrapolating future experimental results. It will guide and be validated by light exotic nuclei studies at FRIB and other exotic beams facilities. Finally, scientists will understand how ^{12}C and ^{16}O, elements critical for life, are produced in nature.

Weak Nuclear Structure—Nuclei as Laboratories for Neutrino Physics

Basic Scientific and Computational Challenges

The neutrino is one of the most elusive particles in the universe and yet one of the most influential. The mass and interaction of the neutrino with other matter are less than a millionth of an electron's—yet neutrinos power spectacular core-collapse supernovae that seed the universe with heavy elements. The fact that neutrinos even *have* a mass is one of the great discoveries of the past 10 years. If the neutrino is its own antiparticle—a so-called Majorana particle—then physics beyond the current standard model of elementary particles must be invoked with consequences impinging upon the matter and antimatter imbalance in the early universe. While the $^{12}C(\alpha,\gamma)^{16}O$ reactions reveal how life can exist, a Majorana neutrino may reveal how matter itself came to exist. However, stringent upper-limits on the existence of neutrino Majorana mass contributions would force scientists to look to other explanations for the fundamental matter-antimatter asymmetry that is observed in the universe today. For recent reviews of current knowledge of neutrino properties, see Avignone et al. (2008), Camilleri et al. (2008), and Haxton (2008).

The primary venue for discerning the fundamental properties of neutrinos is atomic nuclei. A number of experiments are being planned worldwide to determine their properties, but interpreting the results of those experiments will require reliable calculations of nuclear structure and of the interaction between neutrinos and nuclei. Two broad classes of experiments are relevant here, and because of the difficulty in obtaining constraints needed to calibrate these experiments, both require sophisticated theory to be interpreted. As a check on the calculations, as well as a determination of the systematic uncertainty in the theory, scientists will use competing methods to compute the reaction and decay rates.

The first method consists of long-baseline experiments to measure neutrino flavor oscillations, which are sensitive to the differences in neutrino masses, as well as neutrino flavor-mixing angles. Detectors used in these experiments are based on target nuclei such as carbon and oxygen, and it is crucial to understand the neutrino-induced response of these nuclei to fully exploit measurements. At lower energies, neutrino cross-sections on these nuclei also play an important role in late-stage stellar evolution, as well as driving gravitational-collapse supernovae and the creation of heavy elements in supernovae. Reliable calculations require accurate treatment of the strong interaction and a realistic representation of the weak interaction currents. At low energies, the neutrinos couple with nuclei predominantly through so-called "allowed" operators, which are simple, easily calibrated, and cross-checked through experimentation. However, at higher energies—including those relevant to the detectors—scientists also need "forbidden" current operators, which are much more difficult to compare directly to the experiment.

The second experimental methodology consists of neutrinoless double-beta decay (0ν ββ decay) measurements (Vogel 2007). These decays can only occur if the neutrino is its own antiparticle; if so, a neutrino can be emitted and reabsorbed within the same nucleus. If these decays do occur, the lifetime is inversely proportional to the mass of the neutrino and the nuclear matrix element. Unlike 2ν ββ decay, which can be largely calibrated by comparison to ordinary beta decay, the operator responsible for the 0ν ββ-decay nuclear matrix element is neither theoretically simple nor easily constrained by other experiments. Among the specific target nuclei are ^{48}Ca, ^{76}Ge, and ^{130}Te.

The fundamental challenge is to create a computer model of the structure of a nucleus, and then compute the nuclear coupling to neutrinos. Starting from fundamental measurements of NN interactions and using rigorous mathematical methods, effective interactions suitable for use on petascale and extreme scale computers, as well as the weak current operators that describe the interactions of neutrinos with nucleons, will be developed. This will be a significant computational project. A general illustration of the computational requirements for these calculations is provided in Figure 11.

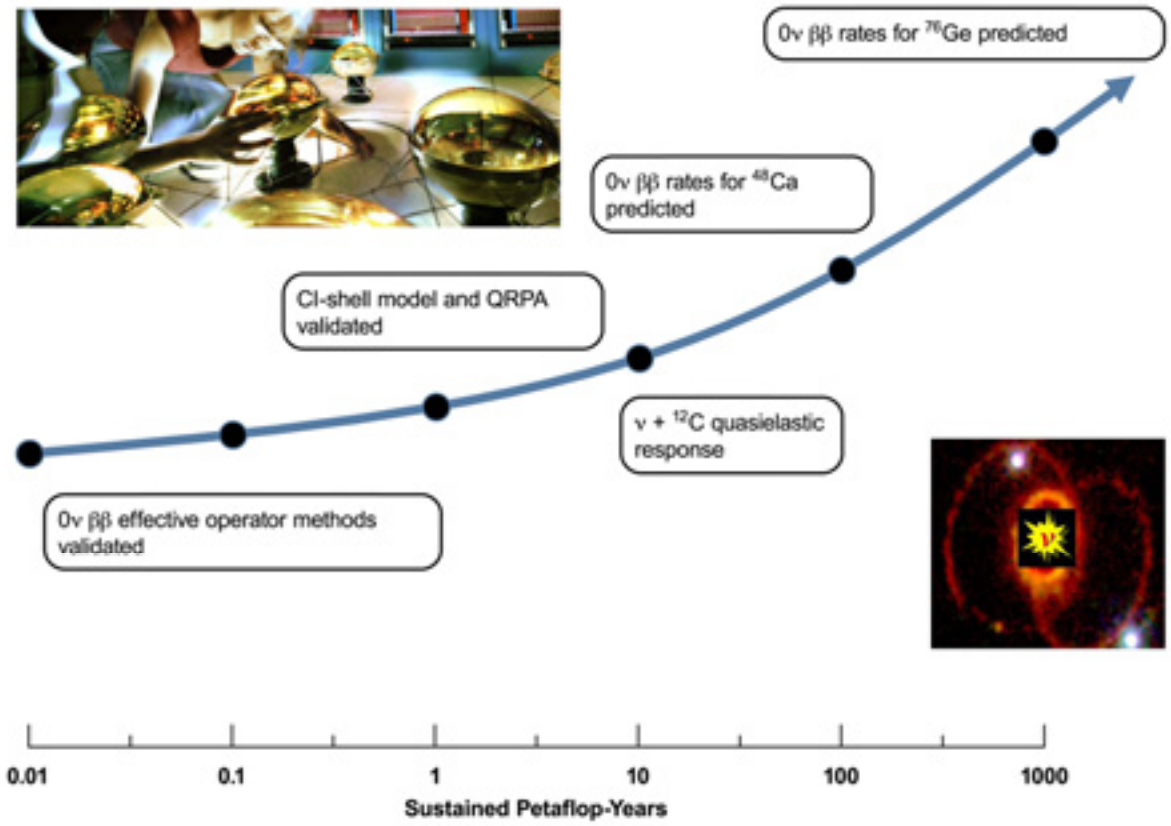

Figure 11. Anticipated highlights for priority research direction "Nuclei as Neutrino Physics Laboratories." Image courtesy of James P. Vary (Iowa State University).

Two main techniques will need to be extended to use extreme scale computing facilities: quantum Monte Carlo (QMC), primarily for the ν-nucleus cross-sections and configuration-interaction shell model (CI-SM) for the 0ν ββ-decay nuclear matrix element and subsequent lifetime. These have complementary strengths and weaknesses. QMC techniques can use "bare" NN and NNN interactions taken directly from

experiment, but QMC cannot yet tackle the heavy nuclei, such as ^{76}Ge or ^{130}Te, relevant to 0ν ββ decay. CI-SM is the technique of choice for detailed spectra and can use arbitrary forms of interactions, not just local potentials—but to fit the problem even on an extreme scale machine, the NN interaction must be renormalized. A third technique, quasi-particle random-phase approximation (QRPA), makes computationally much more modest demands and is thus widely used, but is a more severe approximation.

Each of these techniques faces challenges to be scaled to extreme scale computers. QMC techniques must have actions that are now confined to a single computing core distributed over multiple computing cores. CI-SM will require finding the lowest part of the spectrum of a very large matrix, with dimensions on the order of 1-10 trillion; although the matrix is very sparse, storing the nonzero elements will require petabytes of memory. Furthermore, CI-SM requires vector operations that must communicate across the entire machine. Finally, for CI-SM, scientists must renormalize the experimentally determined NN interaction; this in itself will be a computationally intensive problem because one needs to evaluate the induced NNN interactions in large-basis spaces.

Scientific Outcomes and Impacts

If the neutrino is its own antiparticle, the resulting 0ν ββ-decay lifetime of various nuclei will depend sensitively on the absolute mass of the neutrino. The goal is to compute the 0ν ββ-decay lifetime for nuclei relevant to planned experiments with theoretical uncertainty to 30-50%, cross-checked using competing methods. Accurate estimates of the expected lifetime could affect design of experiments requiring expensive isotopically enriched materials. If a 0ν ββ-decay lifetime is actually measured, these calculations will enable the extraction of the neutrino mass.

Long-baseline oscillation experiments measure the difference between neutrino masses as well as other parameters of the neutrino mass matrix. To correctly interpret the experiments, the ν-nucleus cross-sections will be required to be computed with uncertainties that are less than approximately 20%.

CI-SM can also compute neutrino cross-sections that can provide a cross-check of the QMC and QRPA calculations. As part of this computational project, calculations will be compared using several different methods, usually with the same starting point, from which a systematic uncertainty associated with the calculation can be estimated. One important issue for CI-SM is renormalization, not only of the interaction between nucleons but also between neutrinos and nucleons. Rigorous renormalization methods exist and must be applied consistently to the interaction and the neutrino coupling. Comparisons with results from QMC, where more direct models of the current can be employed, will provide crucial validations.

Currently, significant experimental effort and funds are being invested to answer the above questions, but the experimental results cannot be persuasively evaluated without significant theoretical effort. With extreme scale computing, theoretical studies will provide a basis for reliable interpretation of experiments that explore the properties of neutrinos.

Microscopic Description of Nuclear Fission

Basic Scientific and Computational Challenges

Current understanding of nuclear fission, a fundamental nuclear decay, is still incomplete because of the complexity of the process. Nuclear fission has many societal applications ranging from power generation to national security. In addition, it also plays a role in the synthesis of nuclei in the r-process. Yet, to date, scientists have no microscopic understanding of this complex phenomenon and are unable to make reliable and accurate predictions of fission half-lives, cross-sections, or the distribution of fission products. The ongoing (2009) Scientific Discovery through Advanced Computing Program (SciDAC)-2 Universal Nuclear Energy Density Functional project (Bertsch et al. 2007) and petascale computing resources are opening the way for a comprehensive microscopic description of static properties of atomic nuclei and the fission process.

A promising starting point to obtain a predictive model of nuclear fission is the density functional theory (DFT); see the Nuclear Fission Extreme Scale Computing sidebar. This theory provides the justification for an energy-functional approach to explaining and predicting nuclear structure across the complete table of the nuclides. The accurate nuclear energy functionals currently in use are purely phenomenological and have parameters that are fit to only a subset of nuclear properties. Petascale computing resources and improvements in DFT codes made available through the Universal Nuclear Energy Density Functional project (Bertsch et al. 2007) are opening avenues to the comprehensive microscopic description of complex nuclear phenomena in general, particularly in nuclear fission. Several approaches, each entailing a number of serious computational challenges, can be applied to the description of nuclear fission and will be pursued in this program. The adiabatic approach requires as a first step the determination of the potential energy surface (PES) in a multidimensional space of collective coordinates, which comes from constrained Hartree-Fock-Bogoliubov (HFB) calculations (Warda et al. 2002; Staszczak et al. 2005).

Including all relevant degrees of freedom to obtain a realistic and precise PES is a particularly challenging task. Compounding this issue is the need to evaluate the inertia tensor (Giannoni and Quentin 1980; Warda et al. 2002; Goutte et al. 2005). For this program to succeed, it will be critical to develop suitable algorithms to improve the efficiency of constrained calculations. The imaginary-time HFB (Levit 1980; Arve et al. 1987; Puddu and Negele 1987; Skalski 2008) approach relies on the computation of the full spectrum of dense complex matrices with dimensions that can reach millions. Not all of these matrices are Hermitian. Solving eigenvalue problems of that scale will require an enormous amount of memory, which will create a major bottleneck in the calculations. Heterogeneity in future computer architectures (e.g., use of graphical processing units) will pose another complication. New approaches will therefore be needed to overcome the memory bottleneck in these extreme scale calculations. A general illustration of the computational requirements for these calculations is provided in Figure 12.

Nuclear Fission and Extreme Scale Computing

The United States generates 19% of its electrical power with nuclear power plants; in Europe this figure is 30%. Nuclear power plants produce energy from the fissioning of heavy nuclei. This nuclear fission occurs spontaneously (without an external cause) or when the nuclei are hit by neutrons generated by other fissioning nuclei (induced fission). These are complex processes that are not well understood. This lack of precise knowledge leads to nuclear power plants being built or operated with additional costly safety margins. In principle, increased knowledge of the yields will offer the opportunity to increase, with confidence, the power ratings of existing nuclear reactors and will allow improved design of future reactors.

This situation presents a unique opportunity for nuclear theory to achieve, with the help of extreme scale computers, unprecedented predictive power for both spontaneous and induced fission. The yields under a variety of complex environmental conditions will be investigated in great detail to determine optimum operating conditions. This optimization will account for safety, cost, and efficiency factors at an unprecedented level of accuracy.

The figure below shows an energy surface of the fissioning nucleus ^{258}Fm using one of the currently available approximations. This energy surface determines the fission pathway and the eventual fission energy yield. It is the complexity of this surface—with competing pathways indicated by the dashed lines and superimposed shapes—that makes the problem as challenging as it is, as the yields depend sensitively on subtle differences in the texture of the energy surface. Thus, the predicted dominant fission pathway (the orange dashed line) may not be correct. A more reliable theoretical approach, enabled by extreme scale computers, is needed to make predictions accurate enough to be useful for improved nuclear reactor designs.

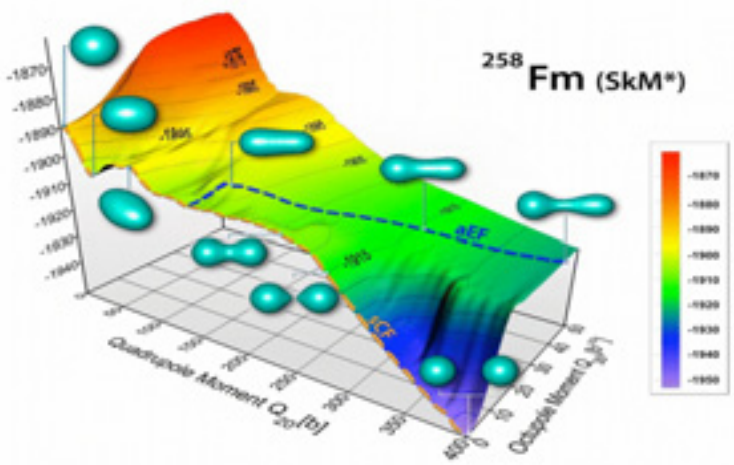

Potential energy surface of ^{258}Fm computed with a standard phenomenological energy-density functional. The fission path follows the line of lowest energy while corrections, such as thermal fluctuations and correlations, lead to alternative nearby paths. The blue figures indicate the shapes taken during fission along different paths. Image courtesy of A. Staszczak, A. Baran, J. Dobaczewski, and W. Nazarewicz (Oak Ridge National Laboratory).

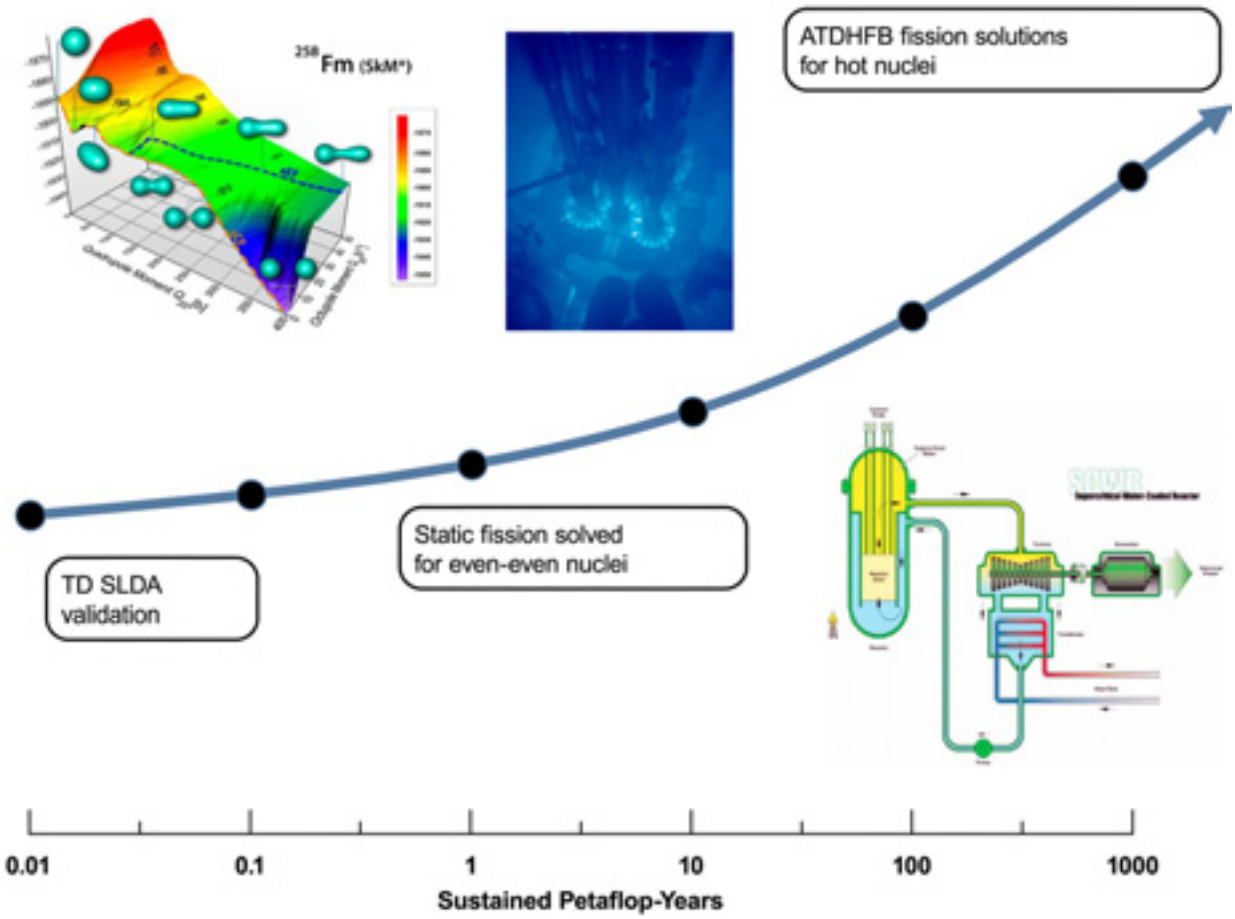

Figure 12. Anticipated highlights for priority research direction "Microscopic Description of Nuclear Fission." Top-left image courtesy of A. Staszczak, A. Baran, J. Dobaczewski, and W. Nazarewicz (Oak Ridge National Laboratory). Remainder of image courtesy of James P. Vary (Iowa State University).

Fission half-lives are extremely sensitive to the details of the underlying PES and the collective mass tensor. This requires extending the current program of energy density functional development to an unprecedented level of precision because phenomenological energy functionals provide essentially a qualitative description. Novel functionals will typically involve 10-30 parameters to be determined through the global minimization of a large number of observables. Constraining effectively each term of the energy functional requires performing symmetry-unrestricted HFB calculations and possibly adopting techniques beyond the mean-field methods. The dimensionality of the problem, combined with the necessity to reach the global minimum, will probably require massive global optimization algorithms. The phenomenon of fission will be investigated with various microscopic approaches. A first step from current capabilities is to follow the adiabatic time-dependent Hartree-Fock-Bogoliubov (ATDHFB) theory.

At least four degrees of freedom—elongation, mass asymmetry, necking, and triaxiality—must be considered. To attain sufficient mesh refinement, it will be necessary to compute of the order of 100,000-plus constrained HFB calculations for every nucleus.

Two nonadiabatic approaches will also be explored. The first is the instanton method, which relies on determining periodic trajectories for the imaginary time HFB equations. Finding the bounce solutions

(periodic instantons) is a difficult numerical challenge. The second approach, applicable in the context of induced fission where the explicit time propagation can be conducted, is a stochastic extension of the time-dependent superfluid local density approximation (TD-SLDA) of DFT. The appeal of this approach, equivalent to the many-body Schrödinger equation, is that two-body and higher correlations become accessible, and dissipation is naturally incorporated into the theoretical description. TD-SLDA has been successfully implemented on current leadership-class super computers, specifically on the Cray XT4 Jaguar at Oak Ridge National Laboratory. A stochastic realization of TD-SLDA will require sufficiently large ensembles of size from thousands to millions of realizations.

Nonadiabatic approaches to spontaneous and induced fission will allow the prediction of the mass and excitation energy distribution of the fission fragments, half-lives, and cross-sections. Beyond the scission point, the emerging fragments start accelerating, and the binding energy of the mother nucleus is converted partially into the kinetic energy of the fragments. At the same time, because strong dissipative processes become increasingly more important, a significant part of the energy is converted into the internal excitation energy of the fragments. The stochastic approach to the time-dependent fission dynamics will allow scientists to calculate these dissipative processes microscopically and predict the nuclear viscosity.

One of the implementation difficulties of stochastic TD-SLDA is the large local memory demand per MPI process and the limited random-access memory/core. Current state-of-the-art calculations prescribe a single MPI process per node so that all the memory in a node is aggregated into a larger, addressable local memory. This approach leaves the other processor cores idle or requires lightweight thread level control within the MPI process to use these cores. Scientists anticipate the need to increase the size of the Hilbert spaces, which will exacerbate this memory-aggregation problem or force the computations out of core—effectively stalling productivity even in the single determinant problems. Programming techniques that go beyond single-node memory aggregation will be refined or developed to satisfy this need. Such developments will also need to include the implicit/explicit use of the extra processor cores.

Scientific Outcomes and Impacts

The computational approach to fission envisioned here, combined with experiments, will provide a predictive framework that may lead to improved nuclear reactor design (AFC 2006). In the area of national security, developing a theoretical description of fission aligns with the goals of the National Nuclear Security Administration Stockpile Stewardship Program, which entails an accurate and complete modeling of the behavior and performance of devices in the nation's aging nuclear weapons stockpile. Improving the accuracy of that description is central to the continuing process of certifying both the safety and the reliability of the stockpile without resumption of nuclear testing and to reduce the threat from nuclear proliferation.

Of all the various nuclear decay processes, nuclear fission—important in the r-process nucleosynthesis, in the modeling of reactions relevant to the advanced fuel cycle for next generation reactors, and in the context of national security—is among the most difficult to tackle. It is a quantum many-body tunneling problem whose typical time-scale changes by orders of magnitude when adding just a few nucleons. The microscopic theory of nuclear fission, rooted in internucleon interactions, still provides a particularly difficult challenge.

The ultimate outcome of the nuclear fission project is a treatment of many-body dynamics that will have wide impacts in nuclear physics and beyond. The computational framework developed in the context of fission will be applied to the variety of phenomena associated with the large amplitude collective motion in nuclei and nuclear matter, molecules, nanostructures, and solids.

Physics of Extreme Neutron-Rich Nuclei and Matter

Basic Scientific and Computational Challenges

Understanding neutron-rich nuclei is vital to discovering the origin of heavy elements (NAP 2003) and defining the properties of neutron-star crusts (Ravenhall et al. 1983). About half of the elements from iron to uranium are produced via successive steps consisting of neutron capture followed by beta decay (the r-process). The structure of neutron-rich nuclei determines the radiative capture cross-sections and beta-decay rates that are critical inputs to r-process nucleosynthesis calculations. The regions around the supposed doubly magic nuclei ^{60}Ca, ^{78}Ni, and ^{132}Sn are of particular interest as they could be waiting points in the r-process. The existence and location of shell closures affect the r-process path as illustrated in Figure 13, where the r-process path is schematically drawn assuming shell closures at the traditional magic numbers. The dynamic and static properties of neutron star crusts determine neutron-star cooling and gravity wave emissions from neutron star mergers.

Unfortunately, present understanding of neutron-rich nuclei is very limited, and extrapolations based on current theoretical models are not reliable. First, the extreme isospin of neutron-rich nuclei magnifies unconstrained properties of the effective nuclear interaction. Second, the proximity of the neutron drip line dramatically increases the number of relevant many-body configurations, including the continuum, and makes accurate computations impossible at the present time. In the coming decade, progress towards the most neutron-rich nuclei will be made with both theory and experiment. The future FRIB at Michigan State University will provide experimental data for selected nuclei along the r-process path. These data will calibrate and validate theoretical methods which, with the advent of exascale computing facilities, will enable accurate theoretical predictions for extremely neutron-rich nuclei (see Figure 13).

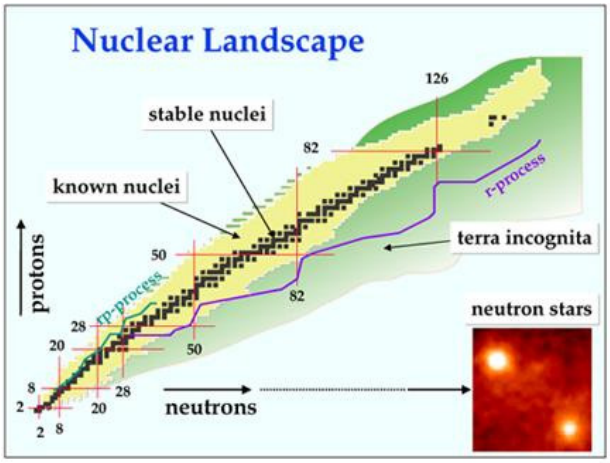

Figure 13. The chart of atomic nuclei displays the speculated r-process path of rapid neutron capture across neutron-rich nuclei. The structure of extremely neutron-rich nuclei is essential input to understand the origin of heavy elements as well as the cooling properties of and the gravity wave emission from neutron star crusts. Image courtesy of James P. Vary (Iowa State University).

The *ab initio* nuclear-structure program aims at building nuclei starting with nucleon degrees of freedom and their mutual interactions. Extending this program to neutron-rich nuclei in the ^{60}Ca, ^{78}Ni, and ^{132}Sn regions and towards the neutron drip lines poses great theoretical and computational challenges. A general picture of the computational requirements for these calculations is illustrated in Figure 14.

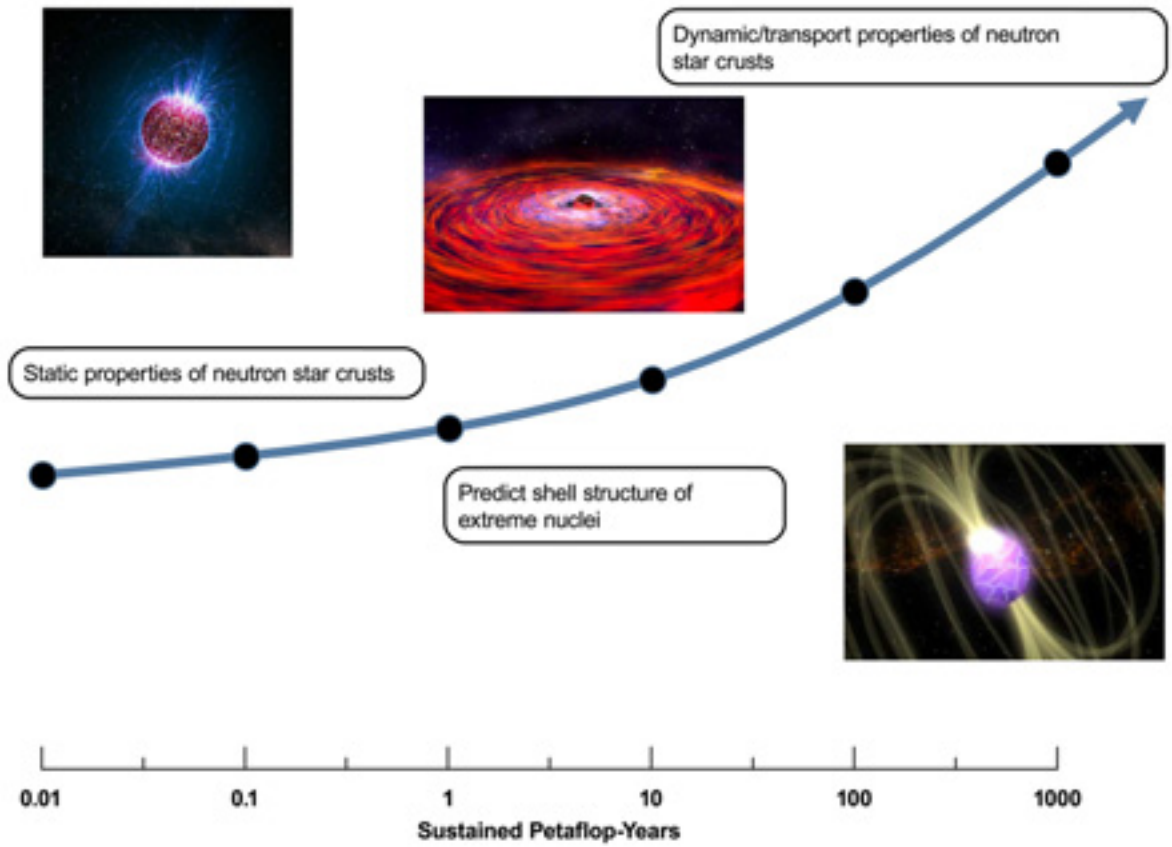

Figure 14. Anticipated highlights for priority research direction "Physics of Extreme Neutron-Rich Nuclei and Matter." Image courtesy of James P. Vary (Iowa State University).

Closed-shell nuclei and their neighbors are of particular interest for both experimental and theoretical research because they form the pillars of understanding and modeling for atomic nuclei.

The effective nuclear Hamiltonian, including the isospin dependence of the effective nuclear two- and many-body forces, is under intense investigation and will become far more precise in the next 3 years. These interactions will be employed with state-of-the-art nuclear-structure tools such as configuration interaction (Lisetskiy et al. 2004), the coupled-cluster method (Hagen et al. 2008), the nuclear density-functional theory (Bertsch et al. 2007), and Monte Carlo techniques (Chang et al. 2004) to calculate the properties of closed-shell nuclei and their neighbors. Of particular interest are the regions around the neutron-rich nuclei ^{78}Ni and ^{132}Sn. These calculations will predict the evolution of shell structure and will explore the drip line and the limits of nucleonic matter. For the understanding of neutron star crusts, the transport properties of systems composed of extremely neutron-rich nuclei and a surrounding neutron gas must be computed.

Calculations of nuclei in the ^{78}Ni region and of static properties of matter in the crust of a neutron star require a facility with tens of petaflop-years of capacity, while computations of nuclei in the ^{132}Sn region and transport properties of crust matter require a facility with hundreds of petaflop-years capacity. Scientists assume the program will be balanced such that investments in computational hardware and software are matched with investments in theory and personnel.

Scientific Outcomes and Impacts

Calculations of nuclei in the ^{78}Ni region and of static properties of matter in the crust of a neutron star require a facility with tens of petaflop-years of capacity, while computations of nuclei in the ^{132}Sn region and transport properties of crust matter require a facility with hundreds of petaflop-years capacity. Scientists assume the program will be balanced such that investments in computational hardware and software are matched with investments in theory and personnel.

These extreme scale computations will allow scientists to determine the limits of nuclear stability—that is, how many neutrons or protons can be bound in a given nucleus. This theoretical effort will have a major impact upon the experimental program to search for these limits at research facilities such as the FRIB. The combination will allow scientists to model some of the most exotic environments in astrophysics, and understand and model the chemical evolution of the universe.

In the crust of neutron stars, neutron-rich nuclei coexist with a surrounding gas of neutrons; the structure and dynamic properties of this unusual matter will be calculated using advanced Monte Carlo methods. In turn, it will be possible to interpret the wealth of astronomical data obtained from visual, x-ray, and gamma-ray telescopes. This will allow scientists to infer details of the nature of these sites and the processes (such as potentially gravitational wave emission) that occur there. The major computational challenge in these efforts is to develop and implement scalable algorithms for the strongly interacting inhomogeneous quantum many-body problem.

An important complement to the work described here will be the experimental program conducted at the FRIB (NRC 2006). The theoretical and computational tools envisioned above will provide an essential framework to interpret FRIB experimental data and will eventually guide the future experimental program. In turn, FRIB data will be essential to verify *ab initio* calculations and calibrate the nuclear many-body Hamiltonian.

Computations of neutron star matter, when combined with observations, will provide information about nucleonic matter at supernuclear densities. The interpretation of observations of isolated, cooling neutron stars require an accurate microscopic understanding of superfluidity and neutrino emission processes in neutron-rich matter. Similarly, observations of gravity waves with the advanced Laser Interferometer Gravitational-Wave Observatory and future detectors will, when combined with a realistic description of the neutron star matter, allow scientists to infer the mass and radius of a neutron star. Combined observations of multiple neutron stars will produce definitive constraints on the equation of the state of the densest matter in the universe.

CONCLUSIONS

Nuclear structure and reaction calculations have consistently made use of available state-of-the-art computers. With steady improvements in nuclear models enabled by computing advances, scientists have reached the ability to make precise predictions of the properties of light nuclei. Extreme scale computing resources will enable such calculations across the periodic table. Besides increasing current understanding of fundamental nuclear physics, these calculations will be of great benefit to other areas including astrophysics, nuclear reactor design, and stockpile stewardship. Progress will rely on continuing a balanced research program, as recent examples illustrate:

- Nuclear physics has benefited from access to the most powerful computers available through grants of time at major computing facilities. The Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is an excellent example.

- Collaboration between nuclear physicists, applied mathematicians, and computer scientists have proven exceptionally fruitful under the Scientific Discovery through Advanced Computing program. There are many formidable obstacles to the efficient use of extreme scale computers.

Table 2 provides an outline of the milestones for the work described in this section. Provided that the computational resources become available for research in nuclear structure and nuclear reactions at the anticipated scales, the forefront research that will be conducted are provided as milestones.

Table 2. Milestones for Nuclear Structure and Nuclear Reactions

Scale	Milestone
>1 Petaflop-year	• Static description of fission for cold even-even nuclei in limited deformation spaces • Time-dependent SLDA • Compute effective transition operators for $0\nu\,\beta\beta$ decay using same method as effective interaction • Side-by-side comparison of $0\nu\,\beta\beta$ decay in CI-shell model and QRPA using same model space/Hamiltonian • Hoyle state in ^{12}C
>20 Petaflop-years	• ATDHFB description of fission in large deformation space • Partial implementation of stochastic TD-SLDA (reduced ensemble) • Moments of quasi-elastic response for ν-^{12}C • Calculations of $0\nu\,\beta\beta$ decay in ^{48}Ca • Tests of current operator; compare methods against experiments • Scattering and capture of α + ^{8}Be • ^{78}Ni structure • Static properties of neutron star crust
>100 Petaflop-years	• ATDHFB description of fission in hot nuclei • Full implementation of stochastic TD-SLDA • Moments of quasi-elastic response for ν-^{16}O • Initial estimates of π-production cross-sections in ν-^{12}C and ν-^{16}O • Converged excitation spectrum of ^{16}O
>1 Exaflop-year	• Complete microscopic description of nuclear fission for odd nuclei such as ^{235}U • Scattering and capture of α + ^{8}Be • ^{132}Sn structure • Calculations of $0\nu\,\beta\beta$ decay in ^{76}Ge • Dynamic/transport properties of neutron star crust

NUCLEAR ASTROPHYSICS

Co-Leads: Anthony Mezzacappa, Oak Ridge National Laboratory
George Fuller, University of California

INTRODUCTION AND CURRENT STATUS

Astrophysics represents the intersection of nearly every important avenue of inquiry and subject area in nuclear physics. Advancing knowledge of the history of the universe requires understanding how quarks, nucleons, and nuclei behave across temperature and density conditions encountered in the very early universe, the interior of stars and supernova explosions, and cosmic ray collisions. Concomitantly, astrophysical considerations have become key tools and goals in efforts to advance many areas of traditional nuclear physics including nuclear structure and reaction theory, ultra dense matter and nuclear matter, and relativistic heavy ion work.

This symbiotic relationship between astrophysics and fundamental nuclear physics is at the center of a remarkable cosmic story. The building blocks (i.e., the periodic table of the elements) necessary for human life have been synthesized over the nearly 14-billion-year history of the baryonic component of the universe, starting with the Big Bang and the formation of the lightest elements. This was followed by the formation of the first stars and the production of heavier elements within those stars, followed by their death in supernovae, and the reprocessing of expelled material in the next generation of stars, and all subsequent star formation, stellar evolution, and stellar death. Certainly, no scientific endeavor is more important than the effort to understand cosmic origins. While scientists generally understand this synthesis, important gaps remain. For example, how do turbulent stellar interiors evolve, and how do stars produce half the elements heavier than iron in their interiors? How do white dwarf stars die in thermonuclear supernovae, and can researchers understand them well enough to probe the nature and evolution of the universe as well as its contents? How do massive stars die in core-collapse supernova explosions, and how do they produce most of the elements between oxygen and iron and the complementary half of the elements heavier than iron? These are all fundamental questions that are germane to our origin and fate, and they remain unanswered.

Other fundamental questions remain unanswered. What is the nature of the “dark energy” that permeates the universe, dominates its mass-energy content, acts repulsively gravitationally, and is responsible for the universe’s accelerated expansion? What are the remaining unknown properties of neutrinos, knowledge of which would provide clues to the origin of all mass in the universe? The universe as a whole, and important events within it, continues to provide a laboratory to explore the physics of the universe’s smallest constituents (i.e., elementary particles, matter at extreme densities and composition, and nuclei). No better example of this exists than the sun. Observations and analysis of neutrinos emitted by the sun played a major role in the discovery of neutrino mass and the determination of other important neutrino properties—both of which revolutionized fundamental neutrino physics, showed that the standard model of elementary particles was incomplete, and provided some of the first clues to a more complete model of elementary particles and unification, a model that scientists still seek.

The foundation laid by computational, experimental, observational, and theoretical efforts to understand the phenomena described here, combined with the promise of extreme scale computing during the next decade, positions researchers to advance the understanding of stars, including the sun, thermonuclear and

core-collapse supernova explosions, and other important phenomena in leaps and not increments. Stars are complicated objects, and only through a demanding series of calculations can the implications of scientific theories be developed and used to understand and interpret observations. True understanding requires the combination of experimentation, the results of theory as worked out through computation, and the new insights that result. In turn, this will lead to answers to some or all of the fundamental questions noted above and provide a more detailed picture of the universe, humankind's origins, and the fate of the universe.

Sun and Other Stars

In the simplest terms, stellar evolution is the theory of self-gravitating, nuclear reactive plasmas. As such, advances in both plasma dynamics and nuclear theory have shaped the understanding of its history. Currently, the biggest obstacle to using astrophysical sites as laboratories for nuclear physics is scientists' rudimentary knowledge of the mixing and transport of mass and angular momentum by turbulence within stars. These mixing processes, driven by convection and the free energy of rotation, play a central role in shaping a star's evolution and mediate the transport of nuclear processed material to the observable surface layers.

The enormous ratio between the hydrodynamic timescale, τ_h, and the nuclear exhaustion timescale, τ_n, during hydrogen, helium, and most phases of carbon burning, renders a full-scale three-dimensional simulation encompassing the entire life span of a star prohibitively expensive, even with extreme scale computing resources. These phases of evolution, where $\tau_n >> \tau_h$, however, offer opportunities to study stellar hydrodynamics during a snapshot of evolution, provide valuable information about mixing processes like convective overshoot mixing, which can be directly incorporated into stellar evolution models. During later stages of evolution—well after the end of the hydrogen burning familiar from the burning of the sun (i.e., beyond carbon burning)—nuclear exhaustion timescales decrease significantly, and scientists find $\tau_n \sim \tau_h$, so that a fully hydrodynamic three-dimensional simulation of the entire burning epoch becomes feasible. Both types of simulations, used in conjunction with stellar evolution codes, are essential for developing a comprehensive understanding of the varied phases of stellar evolution.

Already important strides in stellar evolution modeling are being made through large-scale three-dimensional simulations, providing deep insight into stellar interior physics as diverse as the internal solar rotation profile (Brun and Toomre 2002) and deep mixing within giant stars on the horizontal branch (Eggleton et al. 2006). Complementary to this is the use of stellar evolution model sequences to provide a map between the initial distribution of stellar masses (the initial mass function) and the nucleosynthetic yields (Timmes et al. 1995) that result from both hydrostatic and explosive burning events, including supernovae of Type Ia and II. Large grids of stellar models are used to quantify the uncertainties in predicted abundance patterns caused by the uncertainties in input mixing physics (Woosley and Weaver 1988) and nuclear reaction rates (Tur et al. 2007), thus providing a bridge between large-scale hydrodynamic simulation, stellar evolution theory, and nuclear physics.

Stellar Explosions and Their Remnants: Thermonuclear Supernovae

Type Ia supernovae (SNe Ia) are thermonuclear explosions of massive carbon-oxygen white dwarf stars in binary systems. Understanding the progenitors of these stars is of great interest to learn how they arise. They produce a significant fraction of the heavy elements and most of the iron in the universe. They are also of great importance in cosmology: observations using SNe Ia as "standard candles" revealed the

expansion of the universe is accelerating, and led to the discovery of dark energy (Riess et al. 1998; Perlmutter et al. 1999). Understanding dark energy ranks among the most compelling problems in all of physical science.

Currently, the calibration of SNe Ia as standard candles is entirely empirical. It uses a relation between the peak luminosity of SNe Ia and the rate at which they fade (Phillips 1993). The accuracy of this calibration must be improved from its present value of approximately 15% to better than 1% in order to study quantitatively the behavior of dark energy with redshift (i.e., with the age of the universe) and therefore distinguish among the various explanations of it that have been proposed (Frieman et al. 2008). Providing a better means to calibrate SNe Ia as standard candles is a major goal of three-dimensional simulations of SNe Ias. An independent, reliable method of calibrating SNe Ia would ensure that any evolution of their properties with redshift does not confound their use as standard candles to accurately determine the properties of dark energy.

Comparison of computer simulations and observations led to the conclusion that SNe Ia explosions most likely involve two stages: a buoyancy-driven turbulent nuclear combustion phase that expands the white dwarf star, followed by a detonation phase that incinerates the star and causes it to explode (Hoyle and Fowler 1960; Arnett 1969; Hansen and Wheeler 1969; Nomoto et al. 1976; Khokhlov 1991; Gamezo et al. 2003). However, a full understanding of the explosion mechanism does not yet exist.

This is a consequence of the enormous disparity between the width of the nuclear flame, the detonation wave, the size of the star, the complexity of buoyancy-driven turbulent nuclear combustion, and uncertainty about how the detonation is triggered.

The increasing computational resources available to astrophysicists in recent years have produced deeper insights into buoyancy-driven turbulent nuclear combustions (Khokhlov 1995; Zingale et al. 2005; Zhang et al. 2007). This increase in computational resources has also made possible multiphysics three-dimensional simulations of the explosion (Gamezo et al. 2005; Röpke and Hillebrandt 2005). These simulations led to the discovery of an entirely new explosion mechanism (Plewa et al. 2004; Jordan et al. 2008), demonstrating that such simulations are necessary to capture with sufficient fidelity the physical processes that are involved (see Figure 16).

Stellar Explosions and Their Remnants: Core-Collapse Supernovae

Core-collapse supernovae are among the most violent events in the universe. They mark the death of massive stars (i.e., stars larger than roughly 10 times the size of our own sun) and the birth of neutron stars and stellar-mass black holes. Core-collapse supernovae serve both to synthesize new elements and to disperse elements synthesized in massive stars during their lifetimes. In the end, they are the dominant source of elements between oxygen and iron. They are also one of the only sites in the modern universe where neutrino interactions with matter have macroscopic, dynamic consequences. All of these facts combine to make core-collapse supernovae remarkable cosmic laboratories for nuclear and neutrino physics.

The explosion mechanism of core-collapse supernovae was among the first applications in the history of computational science. The first attempts to form a theory of such supernovae (Hoyle 1946) led to the earliest approaches to simulate these events on computers (Colgate and White 1966). These attempts were soon followed by the first generation of truly multiphysics simulations (Arnett 1966, 1967).

However, the complexity of the problem has meant that, for close to a half century, supernova modelers have struggled to determine the precise nature of the explosion mechanism and to use that understanding to produce quantifiable predictions of the consequences of these events.

In the past few years, as advances in computational power have enabled new levels of simulation, scientists have begun to peel back the layers of feedback-laden uncertainty in their understanding of core-collapse supernovae. Various nucleosynthesis calculations imply that the composition and distribution of the inner nickel-rich ejecta is very sensitive to the details of the explosion mechanism (Fröhlich et al. 2006; Pruet et al. 2006; Kifonidis et al. 2006). Simulations have shown that the neutrino-driven mechanism cannot work for all progenitor masses (Liebendörfer et al. 2001), and that multidimensional effects are advantageous for shock revival in the context of the delayed mechanism (Herant et al. 1994; Burrows et al. 1995; Janka and Müller 1996). Along with the computational discovery of the standing accretion shock instability (Blondin and Mezzacappa 2006), all of these findings lead to the conclusion that multiphysics simulations must be conducted in three spatial dimensions to achieve the requisite physical fidelity (see Figure 16).

Crosscutting Research Direction: Light Curves and Spectra from Thermonuclear and Core-Collapse Supernovae

Although neutrinos and gravitational waves can provide a direct probe of the explosion mechanisms behind some supernovae, scientists generally do not observe the engines directly. Rather, they observe the photons produced in supernova explosions (SNe 1987A being the exception). To understand the engines, the processes in the explosion that create the photons must be understood. In this manner, scientists can tie the observed supernovae to the explosive engines. It is the combination of detailed engine calculations with detailed spectra and light-curve calculations that will allow use of the wealth of astrophysical data to constrain nuclear physics and allow researchers to understand the origin of the elements.

Much of the past work in this field has separated hydrodynamic explosion calculations from radiative transport calculations, using the former as a base for the latter. This can lead to incorrect light-curve estimates and limits what can be learned from spectra. Recently, the first multidimensional, coupled radiation-hydrodynamics spectra and light-curve codes have been developed, proving more sophisticated models are within reach (Woosley et al. 2007; Frey et al. 2009). Realizing the promise of such models requires a quantum advance in computing power.

PRIORITY RESEARCH DIRECTIONS

Sun and Other Stars

Basic Scientific and Computational Challenges

While estimates of nucleosynthetic yields are widely used to infer the origin of the elements that compose the sun and the solar system, as well as to predict the behavior of abundances in the first stars, they depend strongly on the treatment of hydrodynamic mixing in turbulent regions inside stars. Turbulence is a notoriously challenging phenomenon but is ubiquitous in stellar interiors—therefore, a deeper understanding is essential for developing a predictive theory of stellar evolution. Turbulent mixing is a significant problem during the late stages of evolution (post-carbon burning), at which time the nuclear

evolution in the stellar core decouples from the observable surface properties of the star and calibrating the physics of mixing is not possible. Using calibrated mixing rates based on earlier phases of evolution is also not guaranteed to apply during the distinct vigorous core- and shell-burning convection accompanying the late burning stages.

Computationally, modeling stellar turbulence strains the presently available resources because of the enormous range of relevant length and timescales and the variety of physical processes involved. In supernova progenitors and helium-shell flash nucleosynthesis, reactive hydrodynamic flows need to be modeled. This requires a multiple-component fluid description to track the compositional evolution and the associated nuclear energy release, both of which in turn feed back into the dynamics through buoyancy forces. In the case of solar convection, modeling the photosphere involves multiangle, multigroup radiative transfer, which adds considerable extra cost. In addition, magnetic fields are likely to play a nonpassive dynamic role so that a magnetohydrodynamic solution is desirable.

Capturing a high enough Reynolds number (i.e., Re > 1000) is the biggest obstacle that must be overcome to reliably model stellar flows. In a simulation, the effective Reynolds number scales with linear zoning across a domain as $Re \propto N^{4/3}$ for turbulent flow. Therefore, it would be ideal to have approximately 180 zones across each of the relevant scale lengths (large energy-containing eddies) that arise in the flow. In the solar convection zone, the lower convective boundary layer, known as the tachocline, is approximately 10 times narrower than the convection zone depth. Therefore, to resolve this transition layer, approximately ~2,000 ^{3} zones spanning the entire region would be ideal. Fewer zones may be sufficient if an informed choice of nonuniform zoning is used, thereby decreasing the needed zone count by a factor of a few in each dimension. This would lead to an overall reduction of zones by an order of magnitude. Adaptive mesh refinement is not as important as a nonuniform grid for stellar interior modeling because turbulence at high Reynolds numbers is space filling, and scientists are generally interested in studying quasi-equilibrium states. Scalable adaptive mesh refinement methods, however, can provide the underlying computational framework needed to employ a fixed mesh refinement grid on a massively parallel architecture. Additional savings, perhaps as large as $1/\mathcal{M} \approx 100$ in computing time, may be achieved for low-Mach number flows $\mathcal{M} \approx 10^{-2}$, if scalable low-Mach number methods are successfully developed for petascale and exascale platforms.

For each of the three problems highlighted, breakthroughs will be made possible in moving from the petascale to the extreme scale, primarily because of the increased computational volumes and degree of turbulence (i.e., Reynolds number) achievable. In solar convection modeling, petascale resources will afford enough resolution that a turbulent tachocline can be self-consistently incorporated into a global model. Such a simulation would provide a breakthrough in the scientists' ability to understand the heat and angular momentum transport, which is mediated by this boundary layer and determines the differential rotation profile and dynamo action observed in the sun. Extreme scale computational resources would allow for the self-consistent modeling of solar surface granulation within a global circulation model of the sun, providing precision tests of both the simulation techniques and an understanding of the global-scale magnetohydrodynamic activity observed in the active sun.

Three-Dimensional Turbulent Convection in Massive Stars

The art of modeling stars is intimately entwined with a diverse range of astrophysical problems, from stellar population studies and galactic chemical evolution modeling, to the mechanisms underlying supernovae and gamma-ray burst explosions. These diverse topics are bound together, as almost all of this work relies on the ability to accurately model an individual star. Understanding an individual star is deeply connected to the degree at which the turbulent motions that take place deep in the stellar interior are understood. Currently, the exponential growth in computing technology, which is expressed concisely by Moore's law, is beginning a new era of sophistication in the ability to model plasma dynamics within stars. The numerical experiments made possible by the availability of ever-larger pools of computing resources are providing stronger constraints on researchers' theories. These numerical experiments are also inspiring breakthroughs in understanding stellar mixing and combustion processes, the evolution and structure of stars, and the contribution of these stars to the chemical evolution of the universe.

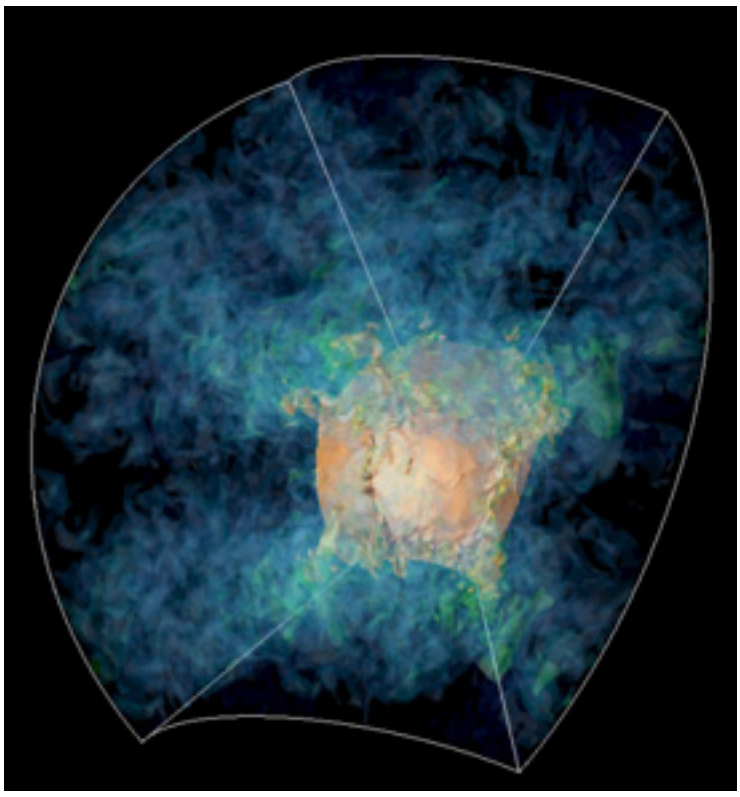

The turbulent convective flow in a core-collapse supernova progenitor star is shown for a large angular domain (120°´ 120°), three-dimensional simulation (Meakin 2008). The mass fraction of ^{32}S is visualized to provide a sense of the topology and the complex, multiscale, turbulent nature of the flow field. Material with a high mass fraction of ^{32}S is being entrained into the turbulent oxygen burning shell from the underlying silicon and sulfur-rich core of a 23 solar mass star. This illustrates a mixing process that is not included in the standard treatment used to model stars today but that has a significant impact on the resulting structure of the massive star at the time of core collapse.

Source: Meakin (2008). Image courtesy of Casey Meakin (University of Arizona).

Developing three-dimensional supernova progenitor models involves a large range in both spatial and temporal scales. While the end state of a massive star depends upon the entire prior evolution of the star since formation, a three-dimensional simulation that begins at core silicon burning and is evolved up to core collapse would provide a significantly improved level of confidence about the state of the iron core at collapse, including the rotational state and the convectively induced perturbations. Such a three-dimensional stellar model would be used directly as an input to core-collapse supernova simulations. Spatially, capturing global asymmetries will require simulating a volume that extends to the outer edge of the carbon burning convection zone (Meakin and Arnett 2006). Thus, in successively larger shells surrounding the core, silicon, oxygen, neon, and carbon burning will need to be included. Angular

momentum transport by wave motions in the stable layers between convection zones (e.g., Talon and Charbonnel 2005) are likely to be important during this epoch and will require a similar computational volume for study. At the petascale, a two-dimensional model encompassing the carbon burning shell could be undertaken and would provide a first-generation multidimensional supernova progenitor model. In a three-dimensional model, the properties of the silicon burning core could be simulated for an hour preceding collapse, thus incorporating realistic features of the vigorous convection. With extreme scale computing resources, a three-dimensional model for the entire silicon burning epoch that incorporates all of the overlying burning shells out to carbon burning would be possible.

The timescale relevant to s-process nucleosynthesis in asymptotic giant branch (AGB) stars is set at a minimum by the period over which helium shell burning convection persists, which is approximately 10 years, while the time period between helium outbursts is approximately 10^3 years. For comparison, the convective turnover time is approximately 3 hours (Herwig et al. 2006) and the Courant time (the hydrodynamic time step limit) is a factor of $f \propto N_{zones}/\mathcal{M} \sim 10^4$ times smaller still, for a Mach number $\mathcal{M} \approx 0.01$ and a modest $N_{zones} \approx 100$ zones spanning the convective shell. Scientists are therefore faced with the problem of evolving the model for an extraordinarily large number of time steps. This temporal problem can be ameliorated by studying snapshots of the quasi-equilibrium turbulent flow. These snapshots guide basic theory to be implemented in stellar evolution codes. This approach requires a three-dimensional simulation spanning only 10 to 100 convective turnovers (Meakin and Arnett 2007). In addition to this temporal challenge is the spatial challenge of capturing the flow in the overlying convective envelope, which extends to very large radii ($r_{env} \approx 3 \times 10^{12}$ cm) compared to the size of the helium burning shell ($r_{shell} \approx 10^9$ cm). While the petascale would allow a first-generation three-dimensional model with sufficient resolution to capture a turbulent convective envelope, the extreme scale would allow for three-dimensional giant star simulations that achieve resolved boundary layer mixing over secular timescales. This is the essential jump from the ability to calculate the bulk to the ability to calculate the surfaces, interfaces, and fine details that yield the observable shape.

The computational developments that would benefit stellar interior modeling include the following: 1) low-Mach number techniques that are scalable to mega-core platforms (e.g., Lin et al. 2006; Almgren et al. 2006), and 2) improvements in reaction network solvers that are informed by reduced quasi-equilibrium-group physics (Hix et al. 2007; Arnett 1996) that can more efficiently treat the complex silicon burning epoch preceding core collapse in massive stars. Treating the solar photosphere self-consistently in a deep convection simulation entails a radiation-hydrodynamics problem, which would benefit from techniques capable of load balancing the multigroup, multiangle radiation transport methods (Nordlund 1982) on a mega-core computing architecture. Finally, a data management challenge is inevitable because of the long integration times necessary to obtain the robust statistics required for studying quasi-steady, turbulent flow. A typical petaflop-scale turbulence simulation with approximately $2{,}000^3$ zones that is sampled 100 times per large eddy turnover for two turnovers would generate approximately 10 petabytes of data if stored at single precision. The total data generated, D, for a turbulence model taking advantage of the available flop rate, F_R, will scale roughly as $D \propto F_R^{3/4}$, so that at the exascale, data volumes should be on the order of a single exabyte.

Scientific Outcomes and Impacts

Stellar evolution, including stellar death through supernovae, answers the question of the origin of the elements in the cosmos. With a firmer knowledge of mixing in stars, the field of stellar evolution theory

and observation will be elevated to that of a precision laboratory for studying the systematics of nuclear matter under extreme conditions, including heavy element nucleosynthesis. The observed solar neutrino flux is already providing important constraints on weak interactions and neutrino oscillation parameters. Extreme scale computing platforms offer exciting new prospects to address several outstanding issues in nuclear astrophysics connected to stellar evolution. Three key areas that will greatly benefit include the following: 1) conducting solar hydrodynamics, 2) performing supernova progenitor modeling, and 3) mixing and nucleosynthesis in giant stars. Figure 15 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

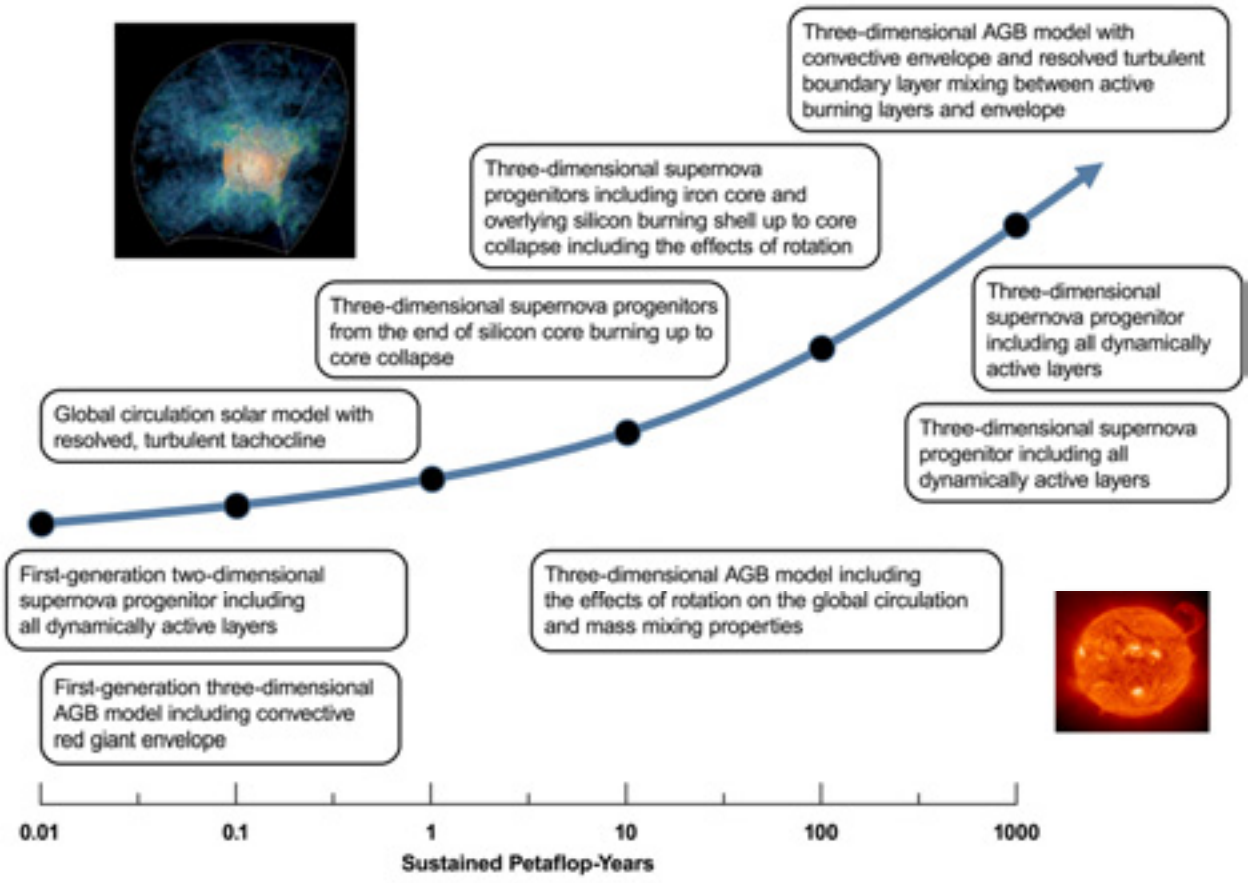

Figure 15. Anticipated highlights for priority research direction "The Sun and Other Stars." Image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

The sun plays a special role as a test of stellar evolution theory because its physical parameters are so well measured. Scientists know its mass, age, radius, and luminosity. Helioseismology has mapped the sound speed to an accuracy of better than 0.5% throughout most of the sun. Solar neutrino spectroscopy has determined the solar core temperature to about 1%. A combination of photospheric and meteoritic measurements constrains solar composition. Scientists can observe the sun's magnetic activity and measure its surface emissions and differential rotations. Modeling this star develops an understanding of the environment on earth and gives scientists an important physics laboratory as the solar neutrino story (Davis 2003) so clearly illustrates.

The one-dimensional standard solar model leaves out many phenomena believed to be important to the sun including convective zone activity, the depletion of light elements in the photosphere, mixing near the radiative/convective zone boundary, and the early convective core—a consequence of out-of-equilibrium carbon burning. The deficiencies of this model are becoming more apparent. Recent three-dimensional modeling of the photosphere—which greatly improved the general consistency of absorption line analyses—has led to reductions in key abundances. In the standard solar model, these abundances must be used throughout the sun, leading to significant changes in sound speeds and a conflict with helioseismology. The differences are most dramatic in the upper radiative zone, where there could be convective overshoot mixing to alter the structure. Alternatively, the standard solar model assumption of homogeneity at zero age main sequence might be incorrect. Ideally, this assumption could be replaced with an explicit three-dimensional calculation of proto-solar formation through collapse of the primordial solar system gas cloud. One speculation connects the photospheric abundance problem with late-stage formation of the planets, which swept out massive quantities of metal from the nebular disk. Therefore, the time is appropriate to bring the level of realism attainable with extreme scale computing to this important astrophysics problem. Aspects of the standard solar model that could be altered in three-dimensional models, including the initial distribution of core metals and the rate of heavy element diffusion, could alter the fluxes of certain neutrino species by up to 20%, limiting the accuracy of the extractable fundamental neutrino parameters. Supernova progenitors evolved in three-dimensional models, which will be possible at the extreme scale, will serve as initial data for both core-collapse and thermonuclear supernova simulations and will provide a significantly improved level of realism over the one-dimensional models that are currently being used. These simulations will provide insight into the complex interplay between convection and weak interactions (the Urca process), which has consequences for the thermal state of a stellar core prior to explosion. The simulations will also address issues related to symmetry breaking by convection. This symmetry breaking in turn seeds instabilities during the supernova event and informs scientists of the rotational state of the stellar core prior to core collapse in massive stars. These tests scenarios proposed to explain gamma-ray burst explosions, which require rapidly rotating cores (MacFadyen and Woosley 1999). In the case of AGB stars, placing knowledge of the mixing occurring in the burning shells and envelopes of these stars on a more solid base enables the production of a predictive model of heavy element formation. These predictive models, used in concert with the copious observational data of the surface abundances in these stars, will provide a powerful laboratory for better understanding element synthesis in the cosmos.

The impact on basic nuclear data is far reaching, as the field of stellar evolution has led to several comprehensive compilations of nuclear data that are used widely in the astronomical community. These data compilations represent active fields of research and have led to standards in the field—such as the rate libraries of Rauscher and Thielemann (2000), which provide a means for assembling experimental and theoretical nuclear physics data from a widely dispersed global effort. These standard libraries of data are easily accessible and enable astrophysicists to explore the broader implications of developments in nuclear theory. For instance, experimentally measured nuclear properties—such as neutron-separation energies and neutron-capture Q values (e.g., see Baruah et al. 2008), and experimentally measured reaction rates, such as $^{12}C(\alpha,\gamma)^{16}O$ (e.g., Assunção et al. 2006),[1] have far-reaching consequences for stellar evolution models and nucleosynthesis (Weaver and Woosley 1993; Tur et al. 2007). Astronomical observations of the abundance patterns across the cosmos—including those in the sun and solar system materials, such as meteorites; low metallicity stars; and giant envelopes—make contact with this input

[1] See panel report titled, "Nuclear Structure and Nuclear Reactions," in this report for further detail on this key reaction.

nuclear physics data and are interpreted explicitly through the scenarios outlined by stellar evolution theory. In addition, observations of specific signatures of nuclear physics, such as the presence of radioactive nuclides in giant envelopes (Cameron 1955; Gallino et al. 1998) and the spectrum of solar neutrinos (Bahcall et al. 2001), are examples of the direct contact that can be made between stellar theory and nuclear physics.

Stellar Explosions and Their Remnants: Thermonuclear Supernovae

Basic Scientific and Computational Challenges

Explosions of SNe Ia involve hundreds of nuclei and thousands of nuclear reactions. These explosions also involve complex hydrodynamic phenomena taking place in degenerate matter and strong gravitational fields (rendering terrestrial experiments of limited utility). Buoyancy-driven turbulent nuclear combustion during the deflagration phase dominates the early part of the explosion and drives an expansion and pulsation of the star. A deflagration-to-detonation transition (DDT) and propagation of the resulting detonation wave through the star has been posited to explain the observed nucleosynthesis and its distribution in space and velocity (Nomoto et al. 1984; Khokhlov 1991; Gamezo et al. 2005). In the alternative gravitationally confined detonation model, fluid flow triggers a detonation that sweeps through the star, producing the observed abundances, spatial distribution, and velocities of the elements.

All of this takes place in approximately 3 s, followed by rapid free expansion of the star at velocities of 10,000 - 25,000 km s^{-1}. These phenomena involve spatial scales from approximately 10^{-3} cm - 10^{9} cm and temporal scales from approximately 10^{-10} s - 10 s, making simulations of SNe Ia a manifestly exascale problem. Advances are needed in both the speed at which the problem can be addressed and the scale (physical size) of the system that can be handled.

Several key physical processes in SNe Ia are not fully understood, and consequently the understanding of the explosion mechanism is uncertain. These physical processes include the smoldering phase, which precedes the explosion phase and is thought to determine the number of points where ignitions occur and their location(s). The buoyancy-driven, turbulent nuclear combustion phase—or deflagration phase, which releases nuclear energy and expands the star—also represents a frontier, as the understanding of reactive turbulence in strong gravity is incomplete. Finally, the origin of the detonation wave that incinerates the star and causes it to explode is uncertain. Whether the physical conditions necessary for a DDT are achieved in the deflagration phase of SNe Ia is unclear. The alternative, in which fluid flow during the deflagration phase triggers the detonation, is not fully understood.

Extreme scale computing resources will produce breakthroughs in the understanding of these physical processes, transforming scientists' ability to simulate SNe Ia. It will enable the qualitative improvement of scientists' understanding of the smoldering phase, thereby reducing the uncertainty in the initial conditions for simulations of the explosion phase. It will make possible studies of buoyancy-driven turbulent nuclear combustions—which include capturing this physical process by simulations that resolve length scales only three to four decades below the largest physical scales and that use a self-similar subgrid model (Khokhlov 1995; Zhang et al. 2007) if needed—that could verify current expectations. If these studies do not verify these expectations, extreme scale computing will determine that the process is more complicated and provide the data needed to construct an appropriate subgrid model. Finally, extreme scale computing will also make possible studies that verify whether buoyancy-driven turbulent nuclear burning in a white dwarf star produces the physical conditions needed for a DDT to occur.

Advances in SNe Ia modeling during the next decade will most likely come from a combination of high-resolution simulations of the key physical processes described above and whole-star simulations of SNe Ia. Sustained petascale computing will enable verification studies of buoyancy-driven turbulent nuclear burning that will dramatically improve the understanding of this key physical process and will make possible whole-star SNe Ia simulations to treat buoyancy-driven turbulent nuclear combustion over a larger range of scales, providing new insights into the energy cascade and instabilities produced by this physical process. With extreme scale computing, it may be possible to attempt first-principle simulations of SNe Ia from ignition through the deflagration phase (i.e., the buoyancy-driven turbulent nuclear burning phase), a difficult problem.

A key component of studies at both the petascale and the extreme scale will be global validation of the models using large numbers of SNe Ia simulations. The need to perform large ensembles of simulations means the average time to perform simulations of adequate resolution will have to be reasonably short to allow for several such simulations to be performed in a given real time. Thus, a careful mix of a few high-fidelity and many low-fidelity simulations will be required. Even so, it means that high-capacity as well as high-capability extreme scale computing platforms will be needed.

Achieving extreme scale computing capabilities for SNe Ia simulations presents several challenges:

- SNe Ia simulation codes need to exhibit strong scaling and run efficiently on platforms with millions of cores and/or that exploit accelerators. Weak scaling will be insufficient because the computational demand scales as the fourth power of the resolution.
- SNe Ia simulations, in common with core-collapse supernova and stellar evolution simulations, require many physical variables per computational cell (e.g., fluid variables, flame variables, nuclear species variables, and radiation transport variables). Thus, the smaller memory per core of future platforms will require the development of new algorithms for efficient domain decomposition and load balancing.
- New parallel input/output (I/O) algorithms need to be developed. These include those that can handle files of many terabytes and beyond, along with mass stores that can accommodate exabytes of data. The turbulent nature of the deflagration phase demands high-temporal resolution in the retained data sets. This leads to the production of remarkable data volumes (many petabytes and possibly exabytes).
- New algorithms for scientific data analysis, including visualization, need to be developed to process petabytes and exabytes of data, along with data archiving techniques that can process exabytes of data and allow for comparative analyses to be performed between huge data sets.

Scientific Outcomes and Impacts

The major scientific outcomes of SNe Ia simulations at the extreme scale will be as follows: 1) ascertaining the explosion mechanism; 2) calibrating SNe Ia as standard candles to an accuracy sufficient to study quantitatively the behavior of dark energy with redshift (i.e., with the age of the universe); and 3) understanding the contribution of SNe Ia to nucleosynthesis. Figure 16 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

Understanding the explosion mechanism will also impact ideas about the interaction of reactive flow and turbulence. The deflagration phase is ultimately a straightforward problem in combustion, trading many

of the complications of terrestrial burning (e.g., geometry of devices, unmixed fuels, soot production, etc.) for far more fundamental ones (e.g., extremely strong gravity, huge Reynolds numbers, and remarkably stiff reaction kinetics). As such, SNe Ia simulations represent unique numerical laboratories in which to explore basic ideas in reactive turbulent flow. The production of realistic SNe Ia simulations will require advances in this basic area.

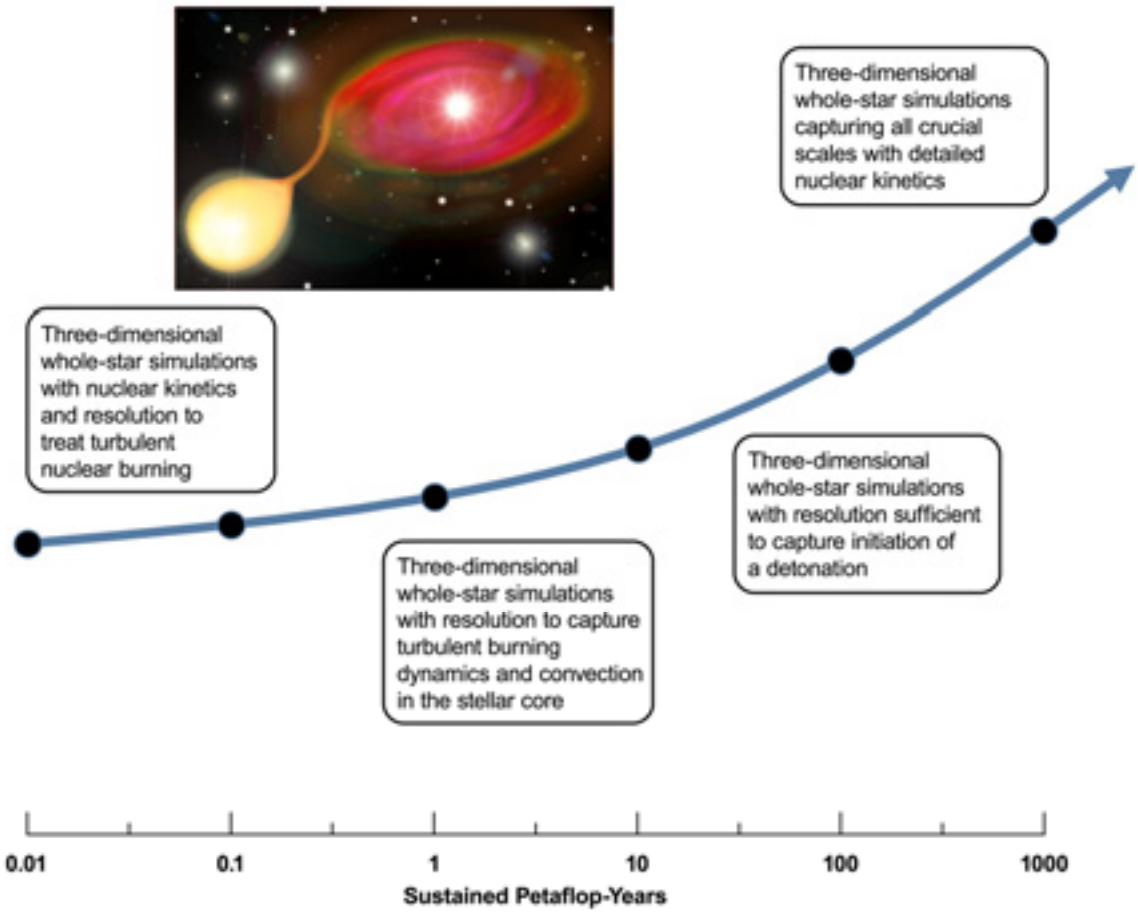

Figure 16. Anticipated highlights for priority research direction "Stellar Explosions and their Remnants: Thermonuclear Supernovae." Upper-left image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

Discovery of Gravitationally Confined Detonation Model of Thermonuclear Supernovae

The ability of large, multiscale, multiphysics simulations to produce breakthroughs in the understanding of thermonuclear supernovae is illustrated by the discovery of the gravitationally confined detonation (GCD) explosion mechanism for Type Ia supernovae (SNe Ia). The development of the FLASH code with adaptive mesh refinement (AMR) and the availability of large computational resources, made possible by the DOE National Nuclear Security Administration Advanced Simulation & Computing Academic Strategic Alliance Program, enabled a team in the FLASH Center at the University of Chicago to perform the first three-dimensional whole-star simulations of the deflagration phase of SNe Ia in 2004. These and subsequent three-dimensional whole-star simulations of SNe Ia showed that, if ignition occurs at one or more points off center (as most scientists think happens), the hot burning bubble that develops following ignition rises rapidly and breaks through the surface of the star, spreads rapidly across the stellar surface, collides at the opposite point, and initiates a detonation (see below figure). The GCD mechanism is the only model to date that detonates without the detonation being inserted by hand, and as such, represents a breakthrough in the SNe Ia field.

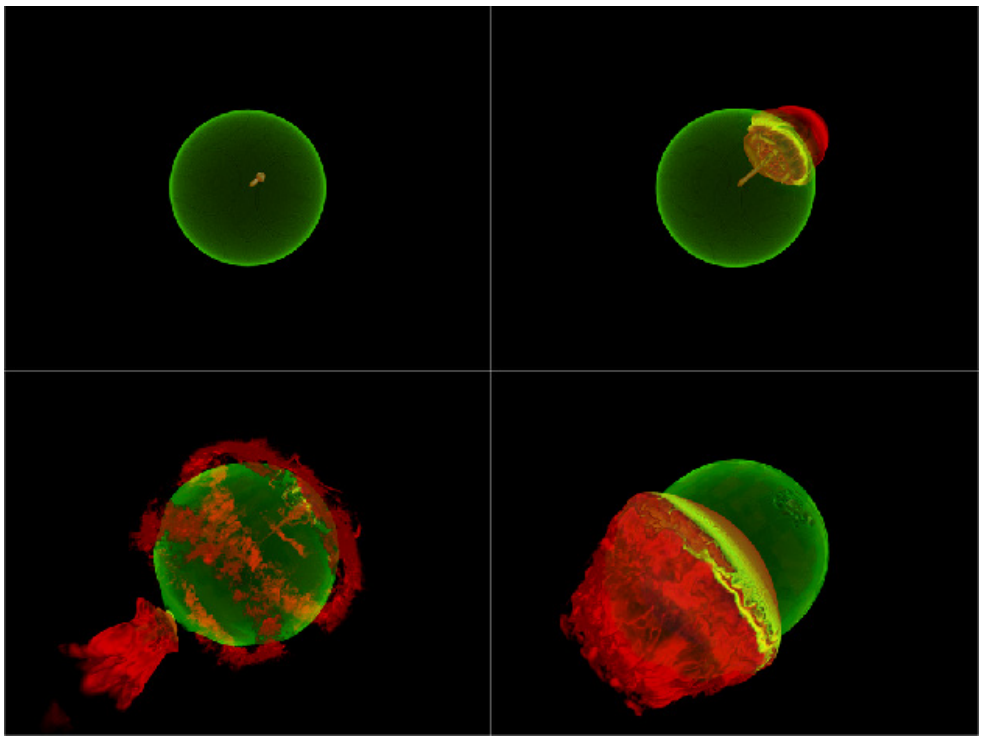

Images in this figure show extremely hot matter (ash or unburned fuel) and the surface of the star at different times for an 8-km resolution simulation of the GCD model starting from initial conditions in which an 18-km radius hot bubble is offset 80 km from the center of the star. The images show volume renderings of the surface of the star and the temperature at (a) 0.5 s, soon after the bubble becomes Rayleigh-Taylor unstable and becomes mushroom-shaped; (b) 1.0 s, as the bubble breaks through the surface of the star; (c) 2.03 s, when the hot ash flowing over the surface of the star has begun to collide; and (d) 2.23 s, as the detonation wave sweeps through the star. Source: Jordan et al. (2008). Image courtesy of Don Lamb (University of Chicago).

An understanding of the explosion mechanism will make possible simulations that can predict correlations among the observed properties of SNe Ia. This will allow them to be better calibrated as standard candles, enabling them to be used to study quantitatively the behavior of dark energy with redshift, and thus to have a strong impact on scientists' understanding of dark energy.

SNe Ia simulations also predict the nucleosynthetic yields for various elements and isotopes, yields that can be tested by observations. These yields are intimately connected with the physical processes that occur during the explosion phase. Consequently, comparisons of nucleosynthetic predictions with observations provide indirect information on these processes, and therefore on the explosion mechanism.

With carbon and oxygen burning being followed by silicon burning and, in the deep interior, an extended period in nuclear statistical equilibrium, SNe Ia simulations are voracious consumers of the nuclear data that govern these burning processes, including binding energies; partition functions; and strong, electromagnetic, and weak interaction reaction rates (e.g., Calder et al. 2007; Seitenzahl et al. 2009). Important reactions, like $^{12}C(\alpha,\gamma)^{16}O$ and triple-alpha burning to form ^{12}C, are the target of ongoing efforts to better measure their reaction rates. Of particular importance are the weak interaction rates for isotopes of iron peak elements, which determine the neutron richness of the simulated ejecta. These continued improvements in the nuclear data improve the nucleosynthetic predictions from SNe Ia simulations, thereby strengthening the constraints on them that are imposed by observations of their ejecta and of solar abundances.

Stellar Explosions and Their Remnants: Core-Collapse Supernovae

Basic Scientific and Computational Challenges

Iron-core collapse and bounce are governed by the interplay of general relativistic gravity with the weak and strong nuclear interactions at extremes of neutron richness and density (e.g., $> 10^{14}$ g/cm^3). The subsequent evolution of the event involves neutrino radiation hydrodynamics and nuclear kinetics among other physical processes. The experimental fact of nonzero neutrino masses means scientists must ultimately solve a macroscopic-scale problem in quantum kinetics as well, directly computing the dynamic, nucleosynthetic, and other observational consequences of flavor oscillations *in situ* as part of fully integrated simulations.

There are profound consequences to this complexity. These simulations make use of a variety of computational algorithms and implementations and stress essentially all facets of a modern, general-purpose computer—the I/O, memory size and latency, processor performance, communication bandwidth and latency, and more—in a manner shared with only a handful of other computational problems. These simulations will stress all facets of a general-purpose supercomputer.

The ability to simulate core-collapse supernovae realistically will depend on the development of discrete representations of the underlying nonlinear partial differential and integro-partial differential equations governing their evolution. This will require efficient and scalable-solution algorithms of the resultant nonlinear algebraic equations, as well as computer codes based on these solution algorithms that can take advantage of the memory and central processing unit capabilities of petascale to extreme scale architectures. Advances in each of these areas will be required, along with considerable work devoted to enhancements of the computational ecosystem surrounding these machines. Core-collapse supernova codes produce prodigious volumes of simulation data over long runs of the codes. Efficiently writing

these data and managing and analyzing them after they are written are as important to producing meaningful science through supernova simulation as is any algorithmic or implementation improvement that might be made for the computational step itself.

Four-Dimensional Core-Collapse Supernovae Simulations

Decades of core-collapse supernovae simulations have led to a challenging proposition for the eventual description of the explosion mechanism and the associated phenomenology. It is now known that three-dimensional hydrodynamics in general relativistic gravity must be coupled to nuclear kinetics capable of accurate estimation of energy release and compositional changes, and to spectral neutrino transport of all flavors of neutrinos (requiring a fourth, neutrino energy, dimension) to reliably determine the nature of the explosion mechanism. Moreover, the inclusion of neutrino masses and oscillations, as well as magnetic fields—both of which may be important effects—make the problem even more difficult. All of these effects, familiar in some sense from the earlier life of the star, operate on very short (millisecond) timescales and at extremes in density (as high as three to four times nuclear matter density) and neutron richness (the ratio of protons to neutrons in the matter can be several times smaller than what is accessible in terrestrial experiments). The rich interplay of all these physical phenomena in core-collapse supernovae ultimately leads to the realization that only through simulation will scientists be able to fully understand these massive stellar explosions. The relative complexity of the germane physics also means the requisite physical fidelity for these simulations will only be realized at the extreme scale.

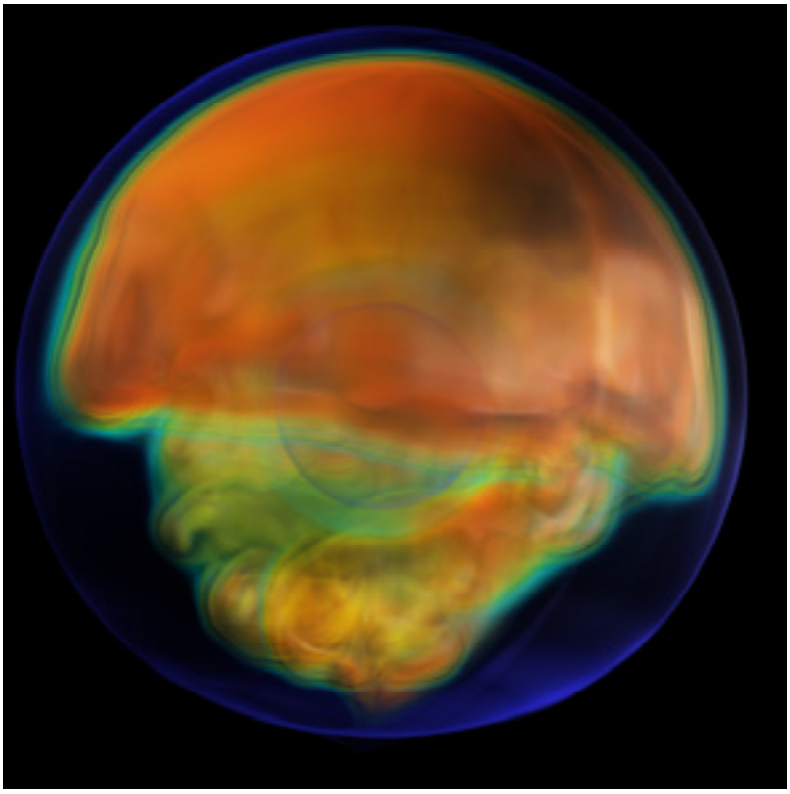

Source: Messer et al. (2009). Image courtesy of Bronson Messer (Oak Ridge National Laboratory).

Rendering of the matter entropy in the core of a massive star at approximately 100 ms following core bounce. Multidimensional fluid effects, coupled to heating and cooling via neutrino emission and absorption, have served to form large, asymmetric structures in the flow. Much of this complicated flow pattern occurs between the surface of the nascent neutron star at the center of the star and the position of the original supernova shock (demarcated by the faint blue circle near the edge of the figure). These first four-dimensional simulations have already consumed tens of millions of central processing unit hours on leadership computing platforms and will require tens of millions of additional central processing unit hours to follow the evolution to the requisite one second of physical time.

Using current petascale platforms and their immediate successors, scientists may be able to determine the general nature of the explosion mechanism itself by performing three-dimensional radiation-magnetohydrodynamics simulations with spectral neutrino transport. As machines capable of peak speeds of 100 petaflops emerge, significant quantitative statements concerning the details of explosive nucleosynthesis in the event and the neutrino emission can be expected. At the extreme scale, scientists will finally be able to determine precisely how supernovae explode by undertaking transformative numerical experiments that incorporate quantum kinetics on macroscopic scales with nuclear physics components realistic enough to accurately predict the isotopic output of these events. These kinds of simulations are utterly unimaginable on current platforms but promise to be accessible at the extreme scale. This is truly applying quantum mechanics, a theory of the smallest things known, to some of the most "macroscopic" bodies in the universe.

Scientific Outcomes and Impacts

The multiphysics nature of core-collapse simulations will require new computational techniques ranging from scalable linear algebra to methods to solve coupled ordinary differential equations. The high number of degrees of freedom at each spatial grid point (e.g., neutrino flavors, energies, and angles, as well as nuclear species) currently represents a large amount of unrealized parallelism in modern supernova codes. Methods to handle these calculations concurrently on multicore platforms and platforms incorporating accelerators of various kinds will likely determine the efficacy of future codes. Figure 17 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

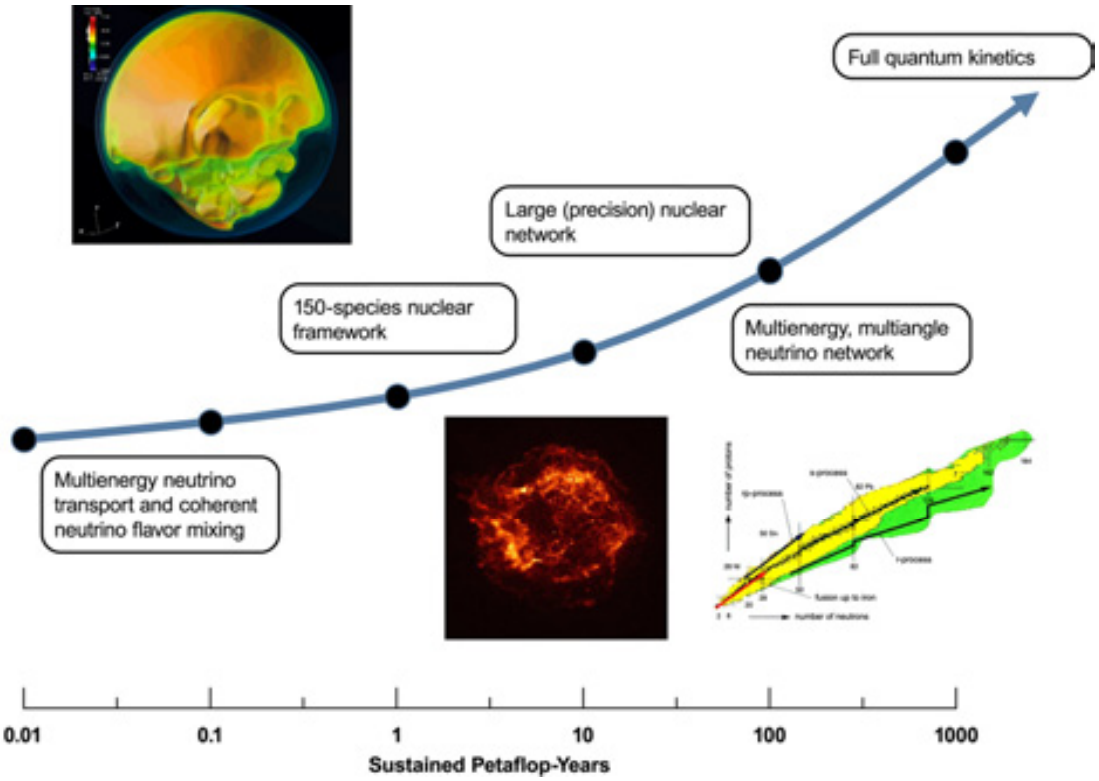

Figure 17. Anticipated highlights for priority research direction "Stellar Explosions and their Remnants: Core-Collapse Supernovae." Bottom-left image courtesy of Chandra X-Ray Observatory and NASA. Remainder of image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

Several of the major questions posed in the 2007 Nuclear Science Long Range Plan (DOE 2007) are germane to core-collapse supernova simulation.

What are the phases of strongly interacting matter, and what roles do they play in the cosmos?

What is the nature of neutron stars and dense nuclear matter?

The nature of dense nuclear matter formed at the center of a supernova explosion provides a unique opportunity to explore the low-temperature, high-density region of the quantum chromodynamics phase diagram. Knowledge obtained from observation and simulation in this region will complement the better-studied, high-temperature (e.g., quark-gluon plasma) regions of the phase diagram, which are presently accessible to terrestrial experiment.

What is the origin of the elements in the cosmos?

One of the most important and distinctive observables from core-collapse supernovae is their pattern of nucleosynthesis. The creation and transmutation of a wide variety of intermediate- and high-mass species in the event is a nonlinear phenomenon. Supernova nucleosynthesis has a dynamic effect on the explosion mechanism, ultimately rendering post-processing of simulation results to be of only qualified utility. The subsequent dissemination of the produced species enriches the interstellar medium, setting the stage for successive generations of star formation and death.

Nuclear physics experiments at the Facility for Rare Isotope Beams, combined with improvements in nuclear theory, will constrain temperature, density, timescales, and neutrino fluxes at the r-process nucleosynthesis site from observations of elemental abundances (RIA Working Group 2006). Simulations of core-collapse supernovae will be the essential ingredients in connecting these experimental measurements to the astrophysical site of the r-process, because a self-consistent determination of all of these conditions can only be achieved through computation at scales beyond those currently possible.

What is the nature of the neutrinos, what are their masses, and how have they shaped the evolution of the universe?

Core-collapse supernovae are, from an energetics point of view, neutrino events. They represent the only instance in the modern universe where neutrino interactions have a discernible, macroscopic effect on the dynamics of baryonic matter. Spectral neutrino transport is required to accurately model the event, and the resulting neutrino templates will be invaluable in interpreting and calibrating detections in terrestrial experiments. Comparing future observations to simulation results will be vital to interpreting those observations and using them to constrain the properties of neutrinos.

Accurate and precise knowledge of the characteristics of neutron-rich matter at high density is a prerequisite for understanding core-collapse supernovae. Precise data for electron-capture processes on progressively larger nuclei is a fundamental need for the simulations, a need which can only be filled by advances in nuclear structure theory. Conversely, core-collapse supernova simulations provide the crucial link in testing these theoretical results, as it is only at the extremes of density and neutron richness realized in these simulations where these predictions are manifest. As nuclei in the collapsing core make the transition from an ensemble of nuclei to nuclear matter, exotic forms of matter are expected (Ravenhall et al. 1983). The details of this transition region are of considerable importance in determining accurate neutrino spectra, again providing a unique link between fundamental theory and physical observables.

In addition, core-collapse supernovae are prodigious sources of gravitational waves (GW) (Ott 2009). Because the signal-to-noise ratio for GW detectors presents a serious complication for detection, the production of useful templates for detectors like the Laser Interferometer Gravitational-Wave Observatory (LIGO) and VIRGO is essential for meaningful data analysis. Furthermore, as nonaxisymmetric oscillations are required for GW production, multidimensional, fully integrated simulations are the only path to producing these signal templates. Therefore, the only path forward to interpreting possible future GW wave detections from core-collapse supernovae relies wholly on simulations providing the requisite context.

CROSSCUTTING RESEARCH DIRECTION: LIGHT CURVES AND SPECTRA FROM THERMONUCLEAR AND CORE-COLLAPSE SUPERNOVAE

Basic Scientific and Computational Challenges

The recent work successfully modeling light curves and spectra with radiation-hydrodynamics codes has opened up new challenges for the light-curve community and a new potential to take full advantage of the growing set of observations. These ties will allow scientists to use supernova observations as laboratories for nuclear physics. The key to using the photons from these explosions to tie the explosion mechanisms, and ultimately nuclear physics, to observations lies in developing coupled, multidimensional, radiation-hydrodynamics code. In this case, special relativity is sufficient for most applications.

The primary computational-physics challenge behind such calculations is similar to the challenge behind core-collapse supernovae: careful coupling of the radiation and hydrodynamics schemes. The opacities, with millions of lines, are much more complicated—an additional constraint on light-curve and spectral calculations. Even assuming the energy levels can be described by a single temperature (local thermodynamic equilibrium), the opacities can only be described by large, two-dimensional tables. However, the situation is even more difficult because many elements are not in local thermodynamic equilibrium, requiring a four-dimensional table and a well-understood radiation field to describe them accurately. The computational difficulty lies in reducing the memory footprint of the multidimensional opacities and working within the memory hierarchy without losing physical accuracy.

Petascale computing is making the testing of the first multidimensional calculations using higher-order transport schemes for the thermal transport and assuming local thermodynamic equilibrium for the opacities possible. However, these calculations are limited to brief durations in time as “proofs of concept.” It will require a 20-petaflop machine to make the first major transformative breakthrough, namely routinely modeling of supernova light curves to use supernova observations to constrain supernova mechanisms and nuclear physics. At this stage, astronomers will be able to make precision estimates for nucleosynthetic yields for some elements and for supernovae at such tight levels that both nuclear physics and explosion mechanisms will, for the first time, feel tight constraints from these observations. The final transformation, to make similar transport calculations but now using opacities out of local thermodynamic equilibrium, will require 5 to 10 years of active research with 100 petaflops to one exaflop of computational power. Clearly, the constraints on nuclear physics and the origin of the elements are strongest if these calculations are coupled with accurate explosion engines and with systems for which scientists have observations of the pre-exploded system (i.e., supernova progenitor stars, which are more common in the latest observational surveys).

Three-Dimensional Spectra and Light Curves from Supernovae: The Ultimate in Validation

Much of what is known about supernovae is derived from observations of their spectra and light curves. Current and proposed surveys (both the National Aeronautics and Space Administration and the National Science Foundation are investing heavily in transient missions ideally suited for supernovae) will increase the rate, spectral resolution, and time resolution of these observations by several orders of magnitude. To connect these data to explosion mechanisms and nucleosynthetic yield calculations, scientists must accurately model the supernova shock as it moves out of a star and into its surrounding medium. Accurate models will require radiation (photon) hydrodynamics calculations using detailed opacities where millions of lines have been accounted for. The modified RAGE (Radiation Adaptive Grid Eulerian) code, an example of a state-of-the-art code, was developed to model supernovae spectra and light curves using the latest opacities assuming local thermodynamic equilibrium for the energy levels of the atoms. Recent results are in the figure below. This code, using implicit Monte Carlo for thermal transport, has been shown to scale well on Roadrunner-type architectures. In addition to shock breakout, this code can, with full opacities, model detailed spectra at all stages of the explosion to directly compare high-resolution spectra to nucleosynthetic yield calculations, ultimately allowing scientists to pinpoint the origin of the elements in the universe. The biggest deficiency in the current calculations is that the opacities are not in local thermodynamic equilibrium. Adding full opacities requires another layer onto the radiation-hydrodynamics calculations. Without such calculations, accurate yield estimates for many elements is impossible. Extreme scale computing will allow the incorporation of out-of-equilibrium opacities, allowing astronomers to finally attain the capability to do precision nuclear physics.

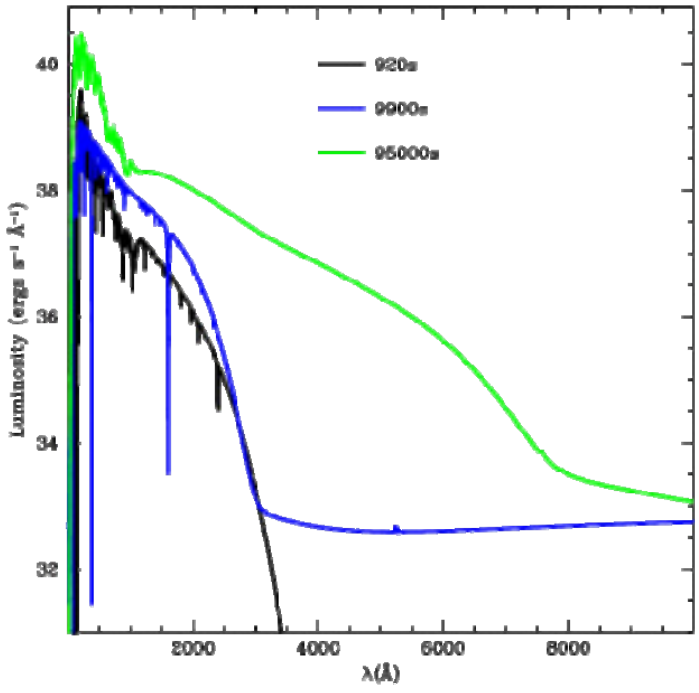

Source: Frey et al. (2009).

Spectra at three different snapshots in time of a Type Ib supernova, showing the evolution of shock breakout (Frey et al. 2009). Shock breakout emission is the burst of high-energy emission that occurs when the supernova shock becomes optically thin (the shock is breaking out of the star or enshrouding wind surrounding the star). This shock breakout has now been observed in a number of cases, providing powerful new probes—especially of the medium directly surrounding the star—into the nature of supernova explosions. With upcoming transient surveys, the quantity and quality of these data will increase by several orders of magnitude.

Scientific Outcomes and Impacts

Light-curve and spectral calculations provide the connection between supernova explosion mechanisms and the rapidly growing wealth of observations from surveys funded by the National Aeronautics and Space Administration (NASA) and the National Science Foundation. As such, the primary goals and impacts of these calculations will be the same as those of the two supernova explosion studies. The development of three-dimensional simulations allows scientists to elevate the entire supernova field to a fully validated science. Some aspects of these will be discussed in the following paragraphs. Any major computational endeavor will also affect computational science.

This effort will lead to radiation-hydrodynamics codes with unprecedented accuracy in the treatment of the opacities. However, the treatment of memory distribution will likely have applications for a much broader scientific community. Already, developing the current codes to run on petascale computing platforms has led to advances in memory use and distribution of information within the memory hierarchy for Monte-Carlo based radiation-hydrodynamics codes (Crawford et al. 2008), leading to a speed-up of this technique over a wide range of architectures. These ideas developed are already spreading to other computational physics challenges.

Constraints on the explosion will allow researchers to test explosion models and the impact that nuclear physics (e.g., variations in the behavior of matter at nuclear densities) has on these explosions. More importantly, however, these calculations will allow astronomers to use observations of cosmic explosions to calculate precise nuclear yields. Nucleosynthesis will move from a quantity integrated over many outbursts to a quantity predicted on a case-by-case basis. Tying the observations to the explosions with these precise calculations could make it possible for astronomers to directly probe nuclear rates and pinpoint the cross sections upon which nuclear yields most crucially depend. With these improved rates, scientists will be in a position to make accurate and predictive estimates for nucleosynthesis.

After moving to precision measurements of the explosions, scientists will most likely discover deficiencies in the current set of nuclear rates used and the description of the behavior of matter at nuclear densities. Astronomers will be able to pinpoint specific areas of study. The nuclear physics community must prepare to answer these specific issues. As such, the nuclear physics requirements identified in the various sections discussing the engines of cosmic explosions will be the same for light curves and spectra.

This schedule matches well with the observational program. Current transient surveys are already increasing the number of observations tenfold (e.g., SuperNova Legacy Survey, Deep Lens Survey, ESSENCE Supernova Survey, and Palomar Transient Factory). In the next 5 to 10 years, another tenfold increase will occur with new surveys (e.g., the Large Synoptic Survey Telescope). With NASA's Swift telescope and the proposed Energetic X-Ray Imaging Survey Telescope (EXIST) mission, astronomers are observing transients in a broader range of wavelengths, providing a fuller picture of the explosion.

This research will also impact nuclear and high-energy density experiments. As crucial rates are determined and erroneous ones are identified, experimental programs will finally be given well-defined cross sections to measure. In addition, these calculations will also pinpoint the most crucial opacities, driving the budding high-energy experimental program studying nonlocal thermodynamic equilibrium opacities.

TOOLS FOR OTHER ASTROPHYSICAL SYSTEMS

The basic building blocks of all astrophysical systems include magnetohydrodynamics; radiation transport (for photons and neutrinos); self-gravity; nuclear kinetics; realistic electromagnetic and weak interaction physics for photons and neutrinos, respectively; and realistic equations of state for nuclear, leptonic (e.g., electrons), and radiation components. Given the ubiquity of these elements across astrophysical phenomena, the computational tools developed by scientists to address the three primary research directions discussed in this panel report can also be used to study other important astrophysical systems that contribute to the origin of the elements and offer other laboratories exploration of new nuclear and neutrino physics.

The Early Universe

One triumph of nuclear physics was the use of Big Bang nucleosynthesis (BBN) to discover the baryon number of the universe. Scientists now know all of the key parameters of BBN except for the lepton number. New physics (e.g., new particles such as sterile neutrinos or decaying dark matter weakly interacting massive particles) can affect light element abundances. This may allow BBN to provide probes not only of new weak-sector neutrino physics but also new mysteries associated with the light element (e.g., ^{6}Li, ^{7}Li, and ^{4}He) primordial abundances and the first stars. However, realizing constraints of this kind in many cases will require a jump to extreme scale computing. For example, calculating the relic densities of sterile neutrinos and the energy spectra of all active and sterile species (i.e., following neutrino flavor evolution in the presence of scattering-induced de-coherence) would require a full solution of the quantum kinetic equations. The development of the solution requires evolving in time and in three spatial dimensions the flavor mixing among three active neutrino species, one or more sterile states, and energy distribution functions for all species. This solution also requires including all elastic and inelastic neutrino scattering contributions and, in the case of conditions near BBN, coupling to all relevant nuclear reactions and abundances.

Novae and X-Ray Bursts

The explosion mechanisms of novae and X-ray bursts (XRBs) are very similar. Both involve systems in which the transfer of typically hydrogen-rich matter onto a degenerate star (white dwarf or neutron star) and thermonuclear ignition of the nuclear fuel under degenerate conditions yield a thermonuclear runaway moderated (at temperatures below about 4×10^{8} K) by the beta-constrained carbon-nitrogen-oxygen cycles. While spherically symmetric numerical simulations (Gehrz et al. 1998; Strohmayer and Bildsten 2006) have reproduced some basic features of these events (e.g., the gross energetics and bolometric light curves), many of the wide range of observations of these objects are not well replicated by such models. Observations have often shown that these events are multidimensional (e.g., many nova nebular remnants have been found to be asymmetric). Given these considerations, clearly only multiphysics and multidimensional simulations can be expected to correctly reproduce the observed behaviors and improve scientists' understanding of these explosive events, of their contributions to nucleosynthesis, and the nature of their progenitor systems.

A significant distinction between the runaways on white dwarfs and those on neutron stars arises from the great discrepancy in their surface gravities. At higher densities, peak temperatures achieved in runaways on neutron stars can reach values approximately 10^{9} K, and the reaction sequences proceed into the proton-rich regions far from stability. Analyses of the light curves of XRBs may yield constraints on

nuclear properties for proton-rich unstable nuclei. Nuclear burning in these events impacts the thermal structure of the crusts of neutron stars and provides constraints on the equation of state of nuclear matter. In contrast, with the exception of nova systems involving more massive white dwarfs and accreting at very low rates, peak temperatures achieved on novae are below approximately 350 million K, and hot carbon-nitrogen-oxygen burning is the dominant reaction. These reaction sequences yield the major nucleosynthesis contributions from novae: the heavy isotopes of carbon (^{13}C), nitrogen (^{15}N), and oxygen (^{17}O). The abundance patterns observed in nova ejecta inform scientists of the burning history and reflect the contributions of novae to galactic nucleosynthesis. It is clear why theoretical and observational studies of novae and XRB systems are considered an important component of the 2007 Nuclear Science Long Range Plan (DOE 2007).

Multidimensional simulations of novae and XRBs to date have been largely limited to two-dimensional models and focused on very specific aspects and regimes. The reason for these limitations becomes clear when the broad range of scales (both temporal and spatial) that are relevant is considered. The nova problem begins with a binary system with a period that is of order hours and dimensions that are of order 10^{10} cm, achieves visual maximum on a timescale of days (at a radius of approximately 10^{12} cm), experiences retreat of the photosphere to a radius of approximately 10^{10} cm on a timescale of order 6 months, and may continue burning for several years. Moreover, at visual maximum, the companion star (mass donor) is orbiting in the envelope of its nova. Glasner et al. (1997) and Kercek et al. (1998) first explored the final stages of the runaways in two-dimensional models, while Alexakis et al. (2004) considered the problem of envelope enrichment via dredge up from the underlying dwarf. If scientists are to understand the effects of convective burning and mixing on nova outbursts and on the composition of nova ejecta, such studies require extension to three-dimensional simulations, inclusion of larger nuclear reaction networks, and coverage of much larger computational domains. For example, expanding the two-dimensional ignition simulations of Glasner et al. (2007) to cover the entire surface of the white dwarf in three-dimensional simulations will require tens of petaflop-years, while further expanding such studies to consider the full range of nuclear species expected to participate in the nova outburst would make this an extreme scale computational problem.

Neutron Star Mergers

Systems composed of two neutron stars or a neutron star and black hole are subjects of some of the most active research in nuclear astrophysics. These mergers offer unique glimpses into the stellar structure at the moment when the stars are obliterated and their components reassembled in a new central object (either a black hole or a larger neutron star) and in a surrounding accretion disk. The gravitational wave signatures of neutron star mergers (NSM) will be able to constrain the stellar equation of state—the inspiral frequency provides an estimate for the neutron star mass and radius while the equation of state stiffness will be observed from the immediate post-merger behavior of the newly formed central object (Rosswog 2007). NSMs have also been postulated as the engines behind short gamma-ray burst phenomena (Mészáros 2006). Mergers of two neutron stars produce thick (about a tenth of a solar mass), opaque accretion disks that are hot and will primarily cool through neutrino radiation. Recent studies (Surman et al. 2008) indicate that the hot inner regions of the disks could be contributors of weak r-process and, to a lesser extent, main r-process elements.

Current simulations model the different NSM stages independently. Full general relativistic simulations usually follow the last orbits of the two objects up to a fraction of a second after the merger; studies of the dynamics of accretion disks usually start with an initial disk surrounding a central object. Finally, the

cooling of the disk by neutrino emission, jet production, and nucleosynthesis is customarily modeled with a semi-stable disk already in place in the presence of a fixed gravitational background. Decadal petaflop systems will likely be used to explore NSM through the last orbits of the binary, the merger, and the early post-merger stage with the formation of the accretion disk and ring-down of the central object. These simulations would cover up to the first dozen seconds after the merger. Following the system's evolution through the late post-merger stage (up to 100 seconds post-merger) will possibly have to wait for platforms in the 100-petaflop range. However, this outline of the future will likely be possible only using basic nuclear networks and approximations to neutrino transport. Accurate neutrino transport and advanced nuclear networks that are central to long-term NSM modeling require extreme scale computing platforms. The advent of next generation supercomputers will coincide with the coming online of the new gravitational wave observatories (NSF's Advanced LIGO 2013 [NSF 2008]; NASA's LISA 2015 [NASA 2007]; European Space Agency's Advanced Virgo 2014 [EGO 2006]). The interpretation of data from these telescopes will be extremely difficult without theoretically produced gravitational wave templates.

BUILDING A COMMON VOCABULARY IN APPLIED MATHEMATICS AND COMPUTER SCIENCE

Extreme scale simulations of turbulent stellar interiors and thermonuclear and core-collapse supernovae lead to an unfolding of the complex evolution in those systems governed by the underlying nonlinear, ordinary differential, partial differential, and partial integro-differential equations that describe stellar core fluid flow and magnetohydrodynamics; self-gravity in both Newtonian and general relativistic systems; radiation transport and radiation hydrodynamics for both photons and neutrinos; and nuclear kinetics. The discretization of these mathematical expressions of physical law leads to nonlinear algebraic equations. Thus, the time evolution of stars and stellar explosions in supernovae is obtained by solving these nonlinear algebraic equations on extreme scale platforms.

For stellar core fluid flow and magnetohydrodynamics, explicit methods have typically been used. Stellar interiors and thermonuclear supernovae are truly turbulent. This requires capturing disparate spatial scales through adaptive mesh refinement and perhaps the implementation of subgrid models, even at the exascale. In some cases, methods for all Mach numbers will be required to traverse disparate timescales. For radiation transport, implicit and Monte-Carlo approaches have been adopted. In the former case, the distributed solution of sparse, structured linear systems of equations with scalable physics-based preconditioners and solvers are needed. These will likely include inner products, consequently requiring scalable global reduction operations, and ultimately will rely heavily on efficient implementations of matrix-vector multiplication on future supercomputing platforms. In the latter case, methods for nonlocal thermodynamic equilibrium transport that fit within future memory hierarchies will need to be developed. Both will become increasingly challenging as the memory per core and memory bandwidth per core decrease. For self-gravity, the elliptic Poisson equation and hyperbolic-elliptic Einstein equations must be solved for Newtonian and general relativistic systems, respectively. This presents well-known challenges of solving such systems scalably on the anticipated megacore platforms. For nuclear kinetics, efficient methods for the solution of nonlinear, stiff, coupled, and large systems of ordinary differential equations are needed. Again, implicit methods have typically been used, given the extreme stiffness of the equations. Note also that nuclear kinetics is a local phenomenon, physically, and hence local to a processor computationally.

For any and all of the above, solution algorithms that are multicore aware will be at the foundation of any success in simulating complex, turbulent systems for scientific discovery. Effective parallel programming models will be needed both across processors and across computing cores within processors. Load balancing of the computation, particularly for simulations that deploy adaptive mesh refinement, will remain important, and simulation fault tolerance as the number of computing cores progresses into the millions will become even more important. Effective debugging tools for simulations that will run on an increasing number of computing cores may be a central challenge. Ultimately, data must be delivered for post-processing. Efficient, collective parallel I/O from millions of computing cores must be established. Data-management approaches for geographically distributed teams and data analysis algorithms must be developed for what will ultimately be petabytes of data per simulation delivered over the course of days to months. Discovery-enabling visualization of multivariate (scalar, vector, and tensor), multidimensional (as high as six-dimensional), petascale data must be developed. Finally, scientific workflows must be made efficient and automated given the expected significant simulation runtimes and the overall daunting data management and visualization challenges.

CONCLUSIONS

Table 3 provides an outline of the milestones for the work described in this section. Provided that the computational resources become available for research in nuclear astrophysics at the anticipated scales, the forefront research that will be conducted are provided as milestones.

Table 3. Milestones for Nuclear Astrophysics

Scale	Milestone
>1 Petaflop-year	• Stellar evolution — First-generation two-dimensional supernovae progenitors — First-generation three-dimensional AGB model with convective envelope — Global circulation solar model with resolved, turbulent tachocline • Core-collapse supernovae — Multienergy neutrino transport with coherent flavor mixing • Thermonuclear supernovae — Three-dimensional whole-star hydrodynamic simulations with resolution sufficient to capture turbulent burning dynamics — Three-dimensional whole-star photon-transport calculations, with resolution sufficient to capture key compositional structures and shape asymmetries, and numbers of atomic lines sufficient to capture key features of optical light curves
> 20 Petaflop-years	• Stellar evolution — Three-dimensional supernova progenitors from the end of silicon core burning up to core collapse • Core-collapse supernovae — Multienergy neutrino transport with coherent flavor mixing • One-hundred-fifty species nuclear network *in situ* • Thermonuclear supernovae — Three-dimensional whole-star hydrodynamic simulations with resolution sufficient to capture turbulent burning dynamics and able to treat the effects on it of convection in the stellar core — Three-dimensional whole-star photon-transport calculations with spatial resolution sufficient to capture key compositional structures and shape asymmetries, and numbers of atomic and molecular lines sufficient to capture key features of optical and near-infrared light curves

Table 3. (contd)

Scale	Milestone
> 100 Petaflop-years	• Stellar evolution — Three-dimensional supernova progenitors, including iron core and overlying silicon burning shell up to core collapse, including the effects of rotation — Three-dimensional AGB model, including the effects of rotation on the global circulation and mass mixing properties • Core-collapse supernovae — Multienergy, multiangle neutrino transport with coherent flavor mixing • Large, precision nuclear network *in situ* • Thermonuclear supernovae — Three-dimensional whole-star hydrodynamic simulations with resolution sufficient to capture turbulent burning dynamics and able to treat the effects on it of convection in the stellar core, and resolution sufficient to capture initiation of a detonation — Three-dimensional whole-star photon-transport calculations with spatial and angular resolution sufficient to capture crucial compositional structure and spherical asymmetry, and numbers of atomic and molecular lines sufficient to capture key features of optical and near-infrared light curves and spectra
>1 Exaflop-year	• Stellar evolution — Global circulation solar model with tachocline and surface granulation — Three-dimensional supernova progenitor, incorporating all dynamically active layers — Three-dimensional AGB model incorporating convective envelope and resolved turbulent boundary layer mixing between active burning layers and envelope • Core-collapse supernovae — Multienergy, multiangle neutrino transport — Full quantum kinetics for neutrinos • Large, precision nuclear network *in situ* • Thermonuclear supernovae — Three-dimensional whole-star hydrodynamic simulations capturing all crucial scales and physical processes, with detailed nuclear kinetics — Three-dimensional whole-star photon-transport calculations capturing all crucial compositional structures, shape asymmetries, and light curve and spectral features

HOT AND DENSE QUANTUM CHROMODYNAMICS

Co-Leads: Steffen A. Bass, Duke University
Frithjof Karsch, Brookhaven National Laboratory

INTRODUCTION AND CURRENT STATUS

Quantum chromodynamics (QCD), the theory of the strong force, is one of the central building blocks of the standard model of particle physics that describes the interaction among all known elementary particles (see sidebar 1 in panel report "Cold Quantum Chromodynamics and Nuclear Forces"). QCD uniquely specifies the interactions between the quarks and gluons. Under ordinary conditions such as they exist in the universe today, quarks and gluons do not appear directly as free particles, but are confined into protons and neutrons. QCD predicts that only under extreme conditions of high temperature or of high density (or both) do the quarks and gluons become the most relevant degrees of freedom that dictate the properties of matter. Temperatures that greatly exceed one hundred million times those controlling the thermal processes on the surface of the sun, or densities that greatly exceed ten times those inside a large nucleus, correspond to such extreme conditions. QCD predicts that in such hot environments, quarks and gluons will behave almost like free particles. This new form of strongly interacting matter is called the quark-gluon plasma (QGP). QGP existed for a short time in the early universe just after the Big Bang, when matter was still hot and dense. However, after a few microseconds this matter cooled down sufficiently so the thermal conditions no longer allowed for the existence of free quarks and gluons. At this stage, the strongly interacting matter went through a cross-over or a phase transition that is reminiscent of the phase transition that occurs when water vapor condenses into the liquid phase. Only after this transition can ordinary matter made out of protons and neutrons be formed in the cosmos.

Deriving detailed predictions of QCD for the properties of matter at high temperature and density is paramount in shaping the current understanding of nuclear matter in general, as well as for understanding the evolution of the early universe. While the properties of nuclear matter at low temperature and moderate densities are well measured, and those of quark-gluon matter at ultra-high temperatures and densities can be rigorously calculated from QCD, little is known about the intermediate regime where the transition between hadronic matter and the QGP occurs. Exploration of this regime, with the aim of mapping the boundaries between different phases of QCD matter, and determining the properties of QCD matter in this domain, is the goal of a large experimental and theoretical program in the United States and internationally. The progress and future goals of this program are described in detail in the 2007 Nuclear Science Long Range Plan, a report issued by the U.S. Department of Energy Office of Science and the National Science Foundation and prepared by the Nuclear Science Advisory Committee (DOE 2007). Large-scale computation plays a pivotal role in achieving the goals outlined in this plan. Numerical calculations are required to describe strongly interacting matter in a regime where collective many-particle effects play a dominant role, which helps bridge the gap between analytic calculations performed in the well-defined thermodynamic limit, and environments that are produced in the laboratory.

Experiments studying QCD matter in the domain of interest are currently performed at the Relativistic Heavy Ion Collider (RHIC) accelerator located at Brookhaven National Laboratory (BNL), and numerical calculations are performed on leadership-class computers at national laboratories, particularly at the BlueGene/L and Blue-Gene/P computers at Lawrence Livermore National Laboratory and BNL, respectively. New experiments are planned for the next decade at RHIC, at the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN), and at the new European Facility for

Antiproton and Ion Research to study the properties of matter not only at high temperatures, but also at a high net baryon number density.

The primary research goal of the entire experimental and theoretical heavy ion physics program in the United States and worldwide is to provide an answer to one of the central questions raised in the 2007 Nuclear Science Long Range Plan (DOE 2007):

What are the phases of strongly interacting matter, and what role do they play in the cosmos?

A central component of the answer to this question is the clarification of whether different regions of the phase diagram of QCD are indeed separated by well-defined phase transition lines where properties of matter change abruptly (see sidebar, "Quantum Chromodynamics Phase Diagram"). Current knowledge of the phase diagram of strongly interacting matter is—to a large extent—based on model-dependent calculations (Stephanov 2006), and very few aspects of the phase-diagram are known from QCD calculations. In particular, it is unknown if a true-phase transition and a line of first-order phase transitions actually exist at intermediate net baryon number densities. These questions can be addressed through numerical calculations performed within the framework of lattice QCD (LQCD), a discretized version of QCD that is formulated on a four-dimensional grid (lattice) (see sidebar 1 in panel report "Cold Quantum Chromodynamics and Nuclear Forces"). Important steps towards a detailed understanding of the phase diagram of strongly interacting matter have been taken. For vanishing net baryon number density, scientists have a reasonable understanding of the transition from ordinary hadronic-matter to the QGP (DeTar 2008). This change in behavior is rapid but smooth; properties of matter change significantly in a narrow temperature, but the cross-over is not accompanied by any singular behavior in observables. Reaching this level of understanding for the QCD phase diagram at nonvanishing net baryon number density (Schmidt 2008) requires extreme computational resources.

An important challenge in studies of strongly interacting nuclear matter and elementary particles at high temperature is to establish contact between the rigorous LQCD calculations, which compute the equilibrium thermodynamics of this matter, and the properties of the strongly interacting matter created in heavy ion experiments. This requires an understanding of dynamic properties of hot and dense matter; e.g., transport properties (Meyer 2008) and in-medium properties of hadrons (Asakawa et al. 2001; Detmold and Savage 2009), and an intensive, computationally demanding microscopic modeling of the rapidly expanding and cooling matter created in heavy ion experiments (Nonaka and Bass 2007). A crucial aspect of such calculations is to determine how the matter—which is originally created in a state far from equilibrium—equilibrates sufficiently rapidly so that a thermodynamic description of its properties becomes possible at relatively early times. Reaching an understanding of the equilibration process in a heavy ion collision requires the modeling of nonequilibrium processes, such as plasma turbulence in a three-dimensional relativistic fluid (TechQM 2008). This is a computationally challenging calculation, requiring new conceptual insights and formal developments.

The Phases of Quantum Chromodynamics

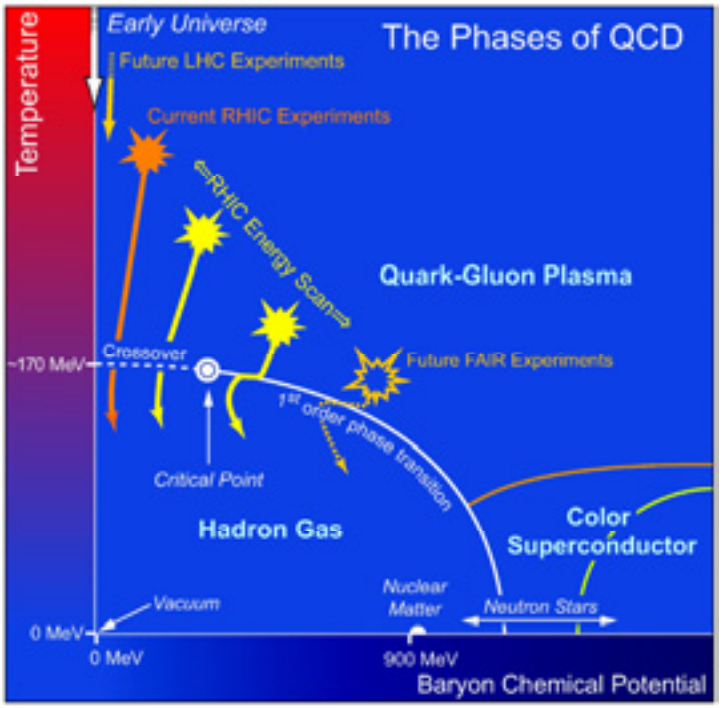

Exploration of the properties of matter as it existed a few microseconds after the Big Bang in the early universe, and as it may still exist today deep inside neutron stars, is subject to extensive experimental and theoretical investigations. Similar to the phase changes that occur when water is heated or compressed (vapor, fluid, solid), nuclear matter (comprised of protons and neutrons) is expected to undergo drastic changes when heated to extremely high temperatures or compressed to high densities.

Strongly interacting matter at high temperatures and/or densities is expected to consist of "deconfined" quarks and gluons, the elementary building blocks of the theory of strong interactions.

Image from *2007 Nuclear Science Long Range Plan: The Frontiers of Nuclear Science* (DOE 2007).

While the transition between the low and high temperature regime is not expected to lead to singularities in thermodynamic quantities at very low net baryon number density, this may be different at larger net baryon number densities. A *critical point* is conjectured to mark the density threshold above which the transition between low- and high-temperature regions is accompanied by discontinuities in baryon number densities, and a latent heat is required to dissolve bound states of quarks—the hadrons—into a new form of matter made of free quarks and gluons. At even higher densities, but low temperature, this matter is predicted to exhibit properties akin to those of superconducting materials.

The theoretical studies of hot and dense matter are performed on state-of-the-art supercomputers. Refining these numerical studies into a precision tool capable of establishing the *phase diagram of strongly interacting matter* requires extreme scale computing resources. Extreme scale computing is also required to perform simulations of dynamical processes describing the experimental results obtained at accelerators such as the Relativistic Heavy Ion Collider and in future experiments at the European facilities Large Hadron Collider and the Facility for Antiproton and Ion Research.

Creating Hot and Dense Matter in the Laboratory

The only known way to create hot and dense quantum chromodynamics matter under controlled conditions in the laboratory—and to investigate its properties—is to collide two heavy nuclei at velocities close to the speed of light with an accelerator. Accelerators currently in operation have the capability of creating the temperatures and densities favorable to the formation of a quark-gluon plasma (QGP). The overarching goal of the experiments performed at these laboratories, is the investigation of the phase diagram of quantum chromodynamics matter, including the deconfined phase, the QGP. The Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (Figure 1) in Long Island, New York, and the accompanying suite of detector systems, were built specifically to observe and study the QGP phase of matter.

Figure 1. The Relativistic Heavy Ion Collider complex at Brookhaven National Laboratory in New York. The complex is comprised of several accelerator facilities joined together to provide beams that are brought into collision in detectors located along the Relativistic Heavy Ion Collider ring. Image courtesy of Brookhaven National Laboratory.

There are four detectors at RHIC: STAR, PHENIX, PHOBOS, and BRAHMS. Two are still active; PHOBOS and BRAHMS completed their operation in 2005 and 2006. Among the two larger detectors, STAR (Figure 2), with its system of time projection chambers covering a large solid angle, is designed for the detection of hadrons; PHENIX is further specialized for detecting rare and electromagnetic interactions.

A typical collision of two gold nuclei, each with momentum of 100 GeV per nucleon, creates a region of QGP matter with a diameter of approximately 10^{-12} cm with a lifetime of approximately 10^{-23} seconds. This QGP fireball then explosively decays into several thousand particles, which have to be tracked and identified by the detectors (see Figure 3). The particle tracking and characterization of the final state of each collision event poses a significant technological challenge to the RHIC experiments, which can record up to several thousand such events per second. The analysis of these events is used to infer the properties of the transient QGP state.

Creating Hot and Dense Matter in the Laboratory (contd)

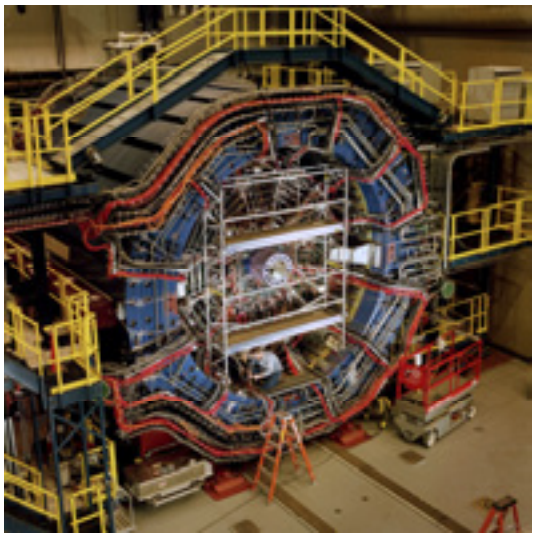

Figure 2. The *Solenoid Tracker at RHIC* (STAR) is a detector designed specifically to track the thousands of particles produced by each heavy ion collision at RHIC. Weighing 1200 tons, and as large as a house, STAR is a massive detector. It is used to search for signatures of the QGP, the form of matter that RHIC was designed to create. Image courtesy of Brookhaven National Laboratory.

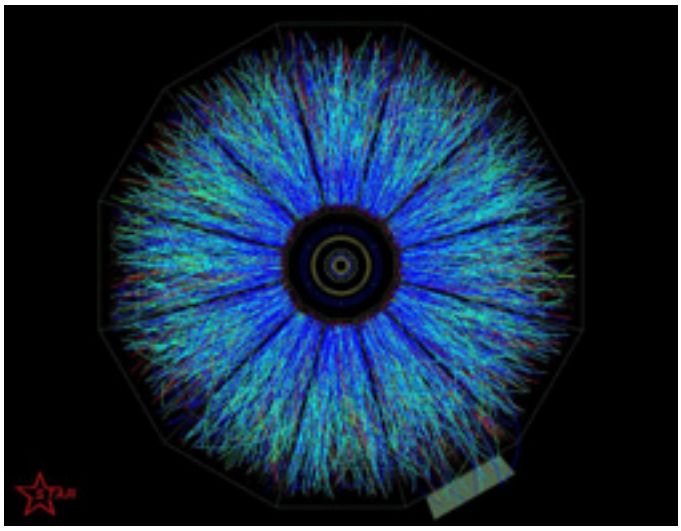

Figure 3. The end view of a collision of two 100-GeV gold beams in the STAR detector at the RHIC at Brookhaven National Laboratory. The beams travel in opposite directions at nearly the speed of light before colliding. Each collision produces thousands of tracks in the detector. Image courtesy of Brookhaven National Laboratory.

RHIC began operation in 2000 and is currently the most powerful heavy ion collider in the world. However, it is expected that the Large Hadron Collider at the European Organization for Nuclear Research (CERN) will provide significantly higher energies after it is fully operational. The planned RHIC-II luminosity upgrade will allow the RHIC and Large Hadron Collider programs to pursue complementary research over the next decade.

The Phase Diagram of Strongly Interacting Matter

The case of vanishing net baryon number density; i.e., the symmetric situation in which the number of particles and antiparticles are identical, plays a special role in the attempt to map out the QCD phase diagram (Stephanov 2006) and to understand phase transitions in strongly interacting matter. Not only does the case of vanishing net baryon number density approximately describe the conditions that existed in the early universe, it is close to conditions that can be studied experimentally in relativistic heavy ion collisions. However, the chiral-limit of QCD, in which the light-quark masses vanish (unphysical values of the quark masses—a theorist's construction) constitutes a theoretically well understood region of the QCD phase diagram where strongly interacting matter is known to undergo a phase transition. In nature, the quarks are not massless, but two of them (the up quark and down quark) have a very small mass compared to the scale of chiral-symmetry breaking. It is expected that the thermodynamics of strongly interacting matter reflects many features of this "nearby" chiral phase transition. Today, scientists have reasonably good constraints on the temperature range over which the transition from hadronic-matter to quark-gluon matter occurs. Nonetheless, to provide useful inputs into the modeling of the dense matter created in heavy ion collisions, a reduction is required in the systematic and statistical uncertainties in current determinations of the transition temperature, of the energy density at the transition point, and of the equation of state over the entire temperature and net baryon number density range covered by current and future heavy ion experiments.

It is important to firmly establish the existence of a second-order phase transition in the chiral limit of QCD. To date, the expected universal scaling properties of various thermodynamic quantities (Karsch and Laermann 1994) have not been reproduced. This raises concerns about the size of the lattice-spacings used in current LQCD calculations, and could indicate that finer lattices are required to firmly establish the continuum limit of QCD in these calculations. More extensive LQCD calculations, including those with improved discretization schemes for the quarks, are required to improve the current state-of-the-art calculations.

Unlike the regime of vanishing net baryon number density, little is known about the QCD phase diagram at nonvanishing net baryon number density through direct numerical calculations. Exploring the structure of the phase diagram at nonvanishing net baryon number density is an outstanding problem that requires numerical approaches quite different from those currently used at zero net baryon number density. As this region of the phase diagram will soon be studied experimentally, it is important to make progress in this area to guide the experimental effort. The challenge is to overcome problems that arise in dealing with extremely high-dimensional integrals that have oscillating integrands. In the context of LQCD calculations at nonzero net baryon number density, this is often called the sign problem. Current numerical approaches that attempt to circumvent this sign problem are promising but require significantly larger computational resources than presently available to reach the level of accuracy required to make quantitative statements about the existence, or nonexistence of a critical point, and a line of first-order phase transitions in the QCD phase diagram. At present, calculations using different approaches to circumvent the sign problem lead to conflicting results on the existence of a critical point (Schmidt 2008; de Forcrand 2009). These conflicting results may be caused by the drastic approximations that have been introduced in current calculations to make them feasible on today's generation of computers.

Equation of State of Strongly Interacting Matter

The equation of state (DeTar 2008)—or more precisely, the temperature dependence of pressure, energy and entropy density that characterize the static, bulk thermodynamic properties of matter—provides basic information on the relevant degrees of freedom that control properties of strongly interacting matter in different regions of the phase diagram. The rapid rise of energy density over a small temperature interval signals the transition from hadronic to quark-gluon matter. Moreover, the value of the energy density in units of the fourth power of the temperature (approximately) counts the number of relevant degrees of freedom at a given value of the temperature. Thus, it is easy to determine that the dominant degrees of freedom at low temperature are hadrons, predominantly pions, while at high temperatures the relevant degrees of freedom are quarks and gluons as suggested by the curves shown in Figure 18. It is more difficult to identify the relevant degrees of freedom in the transition region. Experimental findings at RHIC indicate matter in this regime exhibits properties of a fluid. The existence of various quasi-particle excitations has been postulated to explain the structure of matter in this regime, but it is possible that there are no quasi-particles. Also, the remaining significant deviations from the ideal gas behavior of quarks and gluon at high temperatures raise speculation about QGP properties. Thus, the notion of a quark-gluon fluid has been extensively discussed in the scientific literature, and in view of the experimental findings at RHIC, this fluid has been called a near-perfect fluid. To obtain reliable results on the bulk thermodynamics that allow the verification or falsification of various models of the structure of the high-temperature phase of QCD, it is essential to have precise predictions for bulk thermodynamics over a wide temperature range.

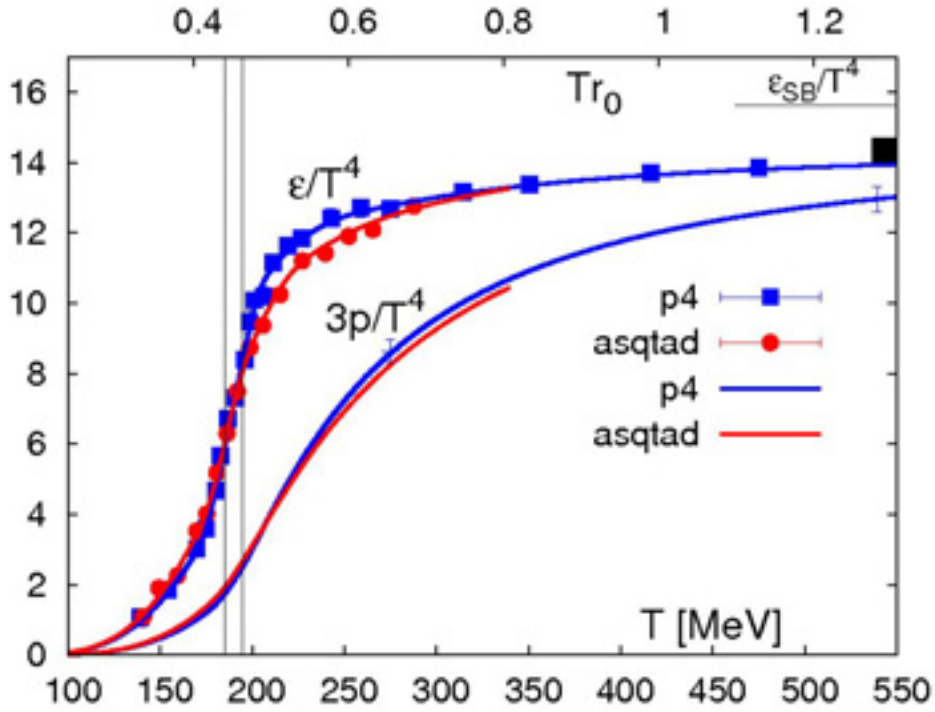

Figure 18. The equation of state of strongly interacting matter calculated with an almost realistic spectrum of quarks for two different discretization schemes (p4 and asqtad). Shown is the energy density and three times the pressure in units of the fourth power of the temperature (Bazavov et al. 2009). The parameter r_0 sets the length-scale in the LQCD calculation. Image courtesy of Frithjof Karsch (Brookhaven National Laboratory).

The Near-Perfect Fluid Nature of the Quark-Gluon Plasma

In April 2005, Brookhaven National Laboratory announced that scientists at the Relativistic Heavy Ion Collider (RHIC) had created the most "ideal fluid" ever observed in nature. "Elliptic flow" is one of the key features at the center of this discovery. In general, heavy nuclei do not collide head-on, but with their centers displaced by an offset called the "impact parameter" (see Figure 1). This leads to the newly created region of highly compressed hot and dense quantum chromodynamics (QCD) matter having the cross-sectional shape of an "almond," with pressure gradients pointing outwards perpendicular to its surface. The shape of the compressed zone, in concert with the resulting pressure gradients, leads to the preferential emission of matter along the impact parameter axis of the overlap zone. This phenomenon is termed "elliptic flow," it is indicative of strong interactions among the involved particles, and can be calculated using relativistic fluid dynamics.

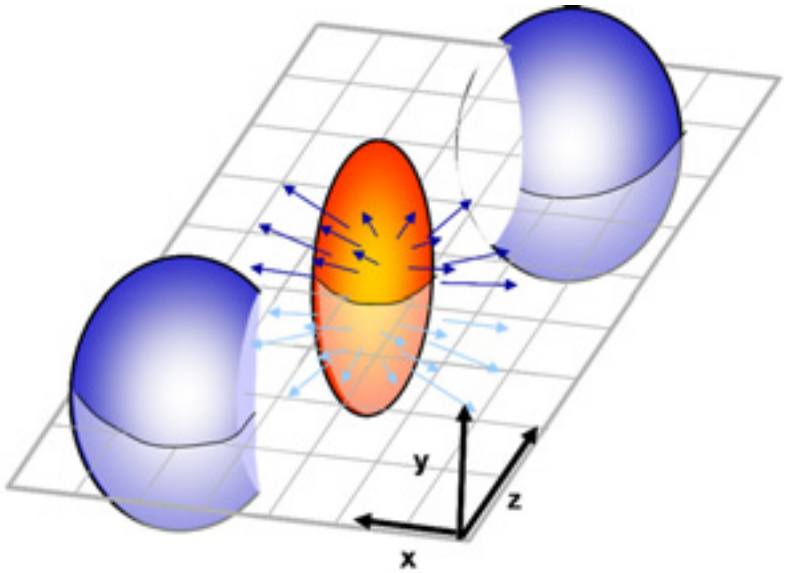

Figure 1. Two nuclei colliding with a nonzero impact parameter create a region of highly compressed QCD matter roughly the shape of an almond. The orientation of the pressure gradients perpendicular to the surface of the almond shape lead to the preferential emission of matter in the xz-plane – this phenomenon is called elliptic flow. Image courtesy of Jerome Lauret.

What makes elliptic flow interesting is that it transforms a transient eccentricity of the matter distribution in coordinate space into a measurable eccentricity of matter in momentum space. Calculations using ideal relativistic fluid dynamics have shown that elliptic flow develops early in the collision, at timescales during which the compressed zone is still in the quark-gluon plasma (QGP) phase. These calculations are in remarkable agreement with RHIC data. The success of ideal relativistic fluid dynamics in describing the elliptic flow observed by scientists conducting the RHIC experiments is consistent with the created QGP having the properties of a near-ideal fluid with a very small shear viscosity to entropy-density ratio, η/s. Figure 2 (next page) shows a viscous relativistic fluid dynamics calculation of the elliptic flow coefficient, v_2, as a function of transverse momentum for several values of η/s (Romatschke and Romatschke 2007). Also shown are the experimentally measured values of v_2 and its uncertainties. First exploratory, but model-dependent, calculations of this ratio in quenched lattice QCD are consistent with the experimental observations (Meyer 2007). Calculations with light dynamical quarks will require extreme scale computing resources. String theory inspired calculations of this ratio yield $\eta/s = 1/4\pi$ in the strong coupling limit of a large class of gauge-theories similar to QCD; unfortunately, such calculations in QCD are not yet possible.

The Near-Perfect Fluid Nature of the Quark-Gluon Plasma (contd)

Most interestingly, however, is that elliptic flow as a manifestation of the near-perfect fluid nature of a system is not restricted to a QGP. The series of pictures in Figure 3 shows elliptic flow occurring in an expanding cloud of ultra-cold lithium atoms being released from an optical trap (O'Hara 2002). Observing the same phenomenon in a system of such different composition and temperature indicates that systems in nature exhibiting elliptic flow may share universal properties.

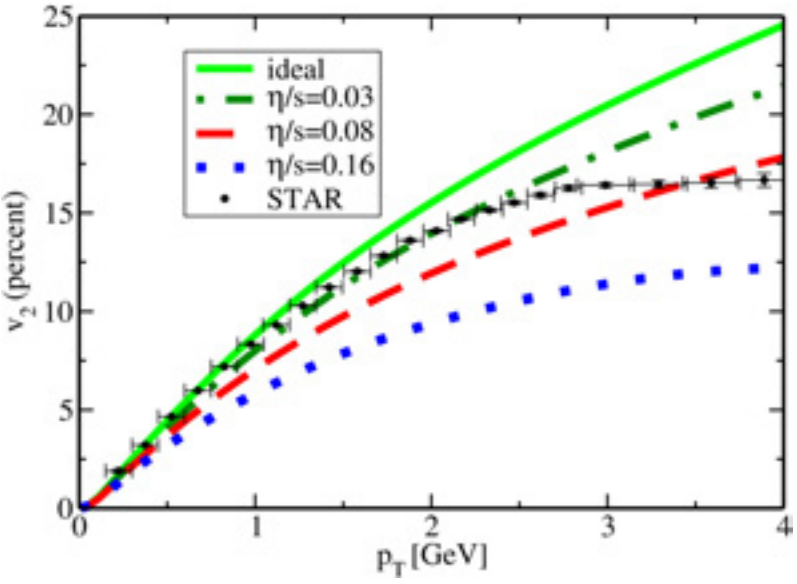

Figure 2. Viscous relativistic fluid dynamics calculations of the elliptic flow coefficient, v2, as a function of transverse momentum for several different values of the ratio of shear viscosity to entropy-density. The experimental values and associated uncertainties obtained by the STAR collaboration are shown by the black points with error-bars. Image courtesy of P. Romatschke (University of Washington).

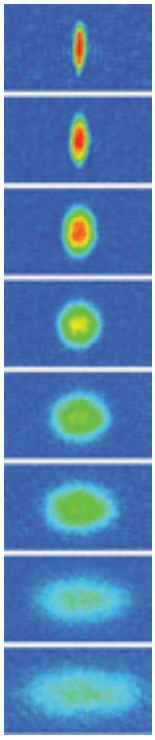

Figure 3. A cloud of ultra-cold fermions (lithium atoms) is released from an optical trap. Due to the initial almond-shape of the trap, the atoms exhibit the same elliptic flow behavior as an expanding QGP formed in ultra-relativistic heavy ion collisions. Source: O'Hara et al. (2002). Image courtesy of John Thomas (Duke University).

To obtain predictions from lattice calculations that are accurate results of QCD, it is necessary to eliminate lattice discretization effects; i.e., perform the continuum limit. To do this in a controlled way requires large-scale numerical calculations. As increasing computing resources have become available, the analysis of the QCD equation of state has been refined. The discretization schemes and algorithms used for numerical calculations have become more sophisticated, leading to a reduction of discretization errors but have become computationally more demanding. Today, studies of the equation of state and basic static properties of the QGP are possible with almost realistic parameters. However, these calculations are extremely time consuming. The most advanced study of the equation of state, which just

has been completed on leadership-class computers at Lawrence Livermore National Laboratory and BNL, required approximately 30 teraflop-years to determine the temperature dependence of the energy density and pressure in the limited temperature range currently accessible at RHIC (Bazavov et al. 2009). Refined calculations at smaller lattice spacings that also cover a larger temperature range, as well as calculations with improved discretization schemes (Jansen 2008) (so-called chiral fermion formulations) are needed in the future to remove the remaining systematic uncertainties in these calculations.

Dynamic Properties of Strongly Interacting Matter

While techniques used to study the static, bulk thermodynamics and the phase diagram of strongly interacting matter are quite advanced, analysis of the dynamic properties remains immature. The temperature dependence of transport coefficients are poorly known, as well as the modification of hadron masses and their widths arising from their interaction with a thermal medium. In-medium modifications of hadron properties are sensitive to the properties of hot and dense matter (Rapp and Wambach 2000), and can be experimentally studied. Transport coefficients are an important ingredient in the modeling of heavy ion collisions. Only recently has a systematic analysis been performed of the experimental data from RHIC using viscous hydrodynamics (Romatschke and Romatschke 2007*)*. This analysis has shown that these data are consistent with a small shear-viscosity-to-entropy.

Information on the transport coefficients and the in-medium properties of hadrons are encoded in the spectral functions that characterize the correlation between external sources put into a thermal medium (Hatsuda 2007). To extract information on spectral functions, high-precision calculations of correlation functions are necessary. A statistical tool, known as the maximum entropy method, can then be used to constrain models of the spectral functions. This approach has many aspects in common with the reconstruction of images from noisy data sets. In the case of QCD, the noisy data are the suitably constructed, numerically evaluated, correlation functions that probe fluctuations of the hot and dense medium. The algorithms used to extract information on transport coefficients and in-medium properties of hadrons from these noisy data are similar to those used in pattern recognition. However, in the context of probing the structure of strongly interacting matter, the main effort is focused on preparing the noisy data. The goal is to reduce the noise level to a point where the filter provided by the maximum entropy method can work efficiently to provide a refined picture of the structure of the QGP. At present, this approach can only be tested in quenched QCD. Even in this context, the information on correlation functions is barely sufficient to allow for a stable reconstruction of spectral functions. To increase the number of time separations that can give information on the correlation functions without reducing the signal-to-noise ratio in the data, numerical calculations are often performed on anisotropic lattices (lattices that have different lattice spacings in the temporal and spatial directions). In the absence of light-quark degrees of freedom, a specific multilevel calculation algorithm is used to reduce the noise level. This approach, however, is not applicable in the presence of dynamical quark degrees of freedom. Much larger computing resources are needed to calculate spectral functions in QCD with all relevant light degrees of freedom.

Microscopic Modeling of Heavy Ion Collisions

Theoretical predictions for thermodynamic properties of strongly interacting matter are based on numerical calculations performed in the framework of LQCD. To explore the properties of hot and dense matter in equilibrium in a realistic system that reflects the properties of QCD correctly, and becomes insensitive to the lattice discretization, requires extraordinary computational resources, as well as the

development of new algorithmic concepts. To establish contact between these first principle studies of equilibrium thermodynamics and conditions met in relativistic heavy ion experiments, an additional theoretical interface is required to model the time evolution and cooling of hot and dense matter. The development of a realistic three-dimensional dynamic modeling tool relies on input from studies of equilibrium thermodynamics and, at the same time, is a computationally highly demanding task. An important goal of this is to obtain a quantitative understanding of the mechanism that leads to the fast apparent equilibration of strongly interacting matter and allows for the observation of thermal effects in heavy ion collisions.

The time evolution of a heavy ion collision at RHIC encompasses several distinct reaction stages, each dominated by very different physical processes (Nonaka and Bass 2007). Figure 19 depicts a schematic view of such a collision: the initial state is comprised of two heavy nuclei (with their wave functions described in terms of elementary quark and gluon degrees of freedom) colliding with each other at approximately 99.9995% speed of light. The system then evolves through a pre-equilibrium stage in which non-Abelian plasma instabilities may drive the system towards local equilibrium. Once nearly equilibrated, the now-formed QGP expands hydrodynamically and, in the process, cools down to the critical temperature of QCD, at which point hadronization occurs. The system of newly formed hadrons continues to interact and expand until freeze out, at which point the individual hadrons cease to interact significantly with each other. Each of these reaction stages has its own effective dynamics and computing challenges that need to be addressed to generate a comprehensive understanding of the time evolution of the collision. Hydrodynamic and particle-based Boltzmann codes are well-suited to describe the latter three reaction stages—these calculations greatly benefit from the recent advent of grid computing. However, the first two reaction stages pose scientific and computational challenges for which petascale, and ultimately, extreme scale computing will be required.

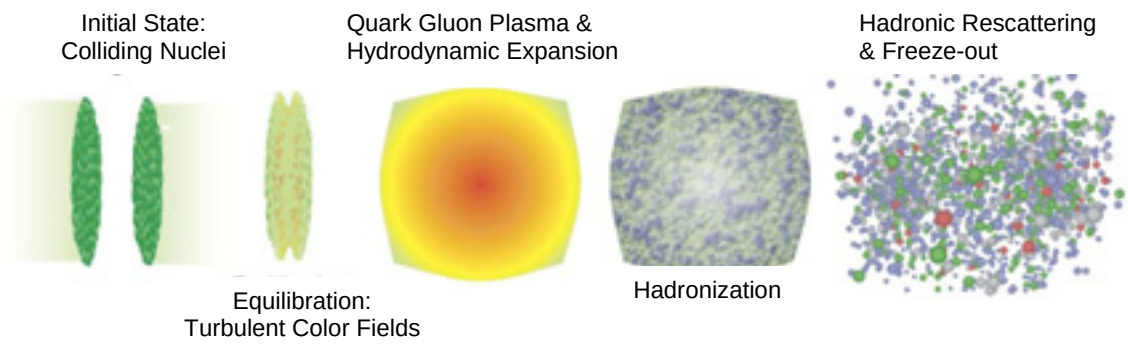

Figure 19. The time evolution of heavy ion collisions at the Relativistic Heavy Ion Collider. The evolution encompasses several distinct reaction stages. Each of these five reaction stages has its own relevant physics and computational challenges that need to be addressed to generate a comprehensive understanding of the time evolution of a collision of relativistic heavy ions. Copyrighted image courtesy of Steffen A. Bass (Duke University).

PRIORITY RESEARCH DIRECTIONS

Precision Calculation of Bulk Thermodynamics

Basic Science and Computational Challenges

Establishing the properties of matter in the vicinity of the chiral phase transition, and characterizing their dependences upon the quark masses and the number of quark flavors will provide fundamental insight into the many remarkable features of QCD. It will enable a study of the interplay between the confinement of quarks and gluons and asymptotic freedom, and a study of the role played by chiral symmetry breaking and topological excitations in generating the masses of the hadrons. Furthermore, establishing the properties of strongly interacting matter in the limit of zero net baryon number density is a prerequisite for any further analysis of the QCD-phase diagram at nonvanishing net baryon number density.

To have complete theoretical control of the thermodynamics of strongly interacting matter in the limit of vanishing net baryon number density, it is necessary to extend the existing calculations of the equation of state and basic static properties of hot and dense matter in several respects: 1) extend scientists current knowledge of the equation of state to higher temperatures; 2) establish better theoretical control over the low temperature regime of the equation of state; and 3) better understand the dependence of thermodynamics on the light-quark masses to be able to explore the phase transition in the chiral limit of massless quarks.

Equation of State

Basic features of the temperature dependence of the energy density and pressure have already been established through LQCD calculations with rather crude approximations to continuum QCD. Calculations on “coarse” (large lattice-spacing) lattices with light-quark masses that are significantly larger than those of nature have shown that a change in the relevant degrees of freedom occurs over a narrow temperature interval (Karsch et al. 2001). However, even with the most current calculations (Bazavov et al. 2009), full control over the structure of the QCD equation of state has yet to be obtained. At high temperatures, contact has not been established with well-defined analytic calculations. At low temperature, the influence of chiral symmetry breaking and its impact upon the hadronic component of the equation of state has not been established. Moreover, the relevant degrees of freedom that control the structure of the equation of state in the transition region have not been determined. Is the restoration of chiral symmetry of any relevance to the QCD transition, or is the copious production of resonances the driving mechanism that leads to deconfinement and a strongly interacting medium of quarks and gluons at high temperature? To answer these questions, LQCD calculations of thermodynamic quantities at lower temperatures must be performed. In addition, lattice discretizations of QCD that respect chiral symmetry, or at least significantly reduce the influence of its explicit breaking due to the finite lattice spacing, are required in the transition region.

Chiral Fermions

To go beyond the current state-of-the art calculations of the QCD equation of state, it is necessary to use improved discretization schemes for the QCD action that respect all of the symmetries of the continuum theory. These discretization schemes have been developed over several years and continue to be

improved through further development. However, these discretization schemes have not been used extensively for numerical studies of QCD to date. This is because they require significantly larger computational resources to perform calculations with sufficiently small statistical uncertainties to allow for a meaningful comparison with the numerical results obtained with non-chiral discretizations.

While improved staggered fermion actions like the highly improved staggered quark (Follana et al. 2007) and stout (Morningstar and Peardon 2004) actions will be used extensively on petaflop computers, truly chiral formulations—such as domain wall and overlap fermion actions (Jansen 2008) — will require extreme scale computing resources in order for a comprehensive study of chiral aspects of the QCD equation of state. These discretized versions of the QCD action provide significantly better control over the chiral properties of QCD, and thus will be important for analyzing the low temperature and transition region of the static, bulk thermodynamic observables, for calculating hadronic screening lengths (Detmold and Savage 2009), and for determining order parameters that characterize the state of matter at high temperatures. Calculations with chiral fermions will enable the analysis of the universal properties of the transition, such as the scaling behavior of the chiral condensate, its susceptibility as well as quark number susceptibilities, and their fluctuations. Further, this work will provide a clarification of the relation between the QCD equation of state and the phenomena of deconfinement and chiral symmetry restoration.

High-Temperature Limit

Properties of strongly interacting matter at temperatures as large as three to four times the transition temperature will be probed experimentally at the LHC at the European Organization for Nuclear Research (CERN), Switzerland. At these temperatures, it may begin to be possible to make contact with perturbative calculations in finite temperature and density QCD (Kajantie et al. 2003; Vuorinen 2003). This will allow for a cross-check between numerical and analytic techniques used in this regime. A reliable numerical calculation of the equation of state and various screening lengths at such high temperatures requires large computational resources as large lattices are needed to control the renormalization of thermodynamic quantities through a proper subtraction of zero temperature observables. This allows for an elimination of otherwise divergent contributions that would prohibit a controlled extrapolation to the continuum limit. Recently developed techniques that minimize the required input from large zero-temperature calculations (Endrodi et al. 2007; Umeda et al. 2009) have the potential to make these calculations less demanding.

Computational Challenge

Calculations with domain wall fermions or overlap fermions require approximately two orders of magnitude more computational resources than calculations performed with staggered fermions. Prospects for the next generation of studies of bulk thermodynamics based on the staggered fermion discretization scheme have been examined in a white paper written in 2007 by the USQCD collaboration (USQCD 2007). This led to the conclusion that a thorough analysis of the equation of state at temperatures below twice the transition temperature will require approximately 100 sustained teraflop-years. Extending such a study to temperatures twice as high will increase the numerical effort by almost an order of magnitude. A thorough study of the QCD equation of state in the transition from low to high temperature needs to be performed with domain wall or overlap fermions. Such calculations require extreme scale computing resources, as shown in Figure 20.

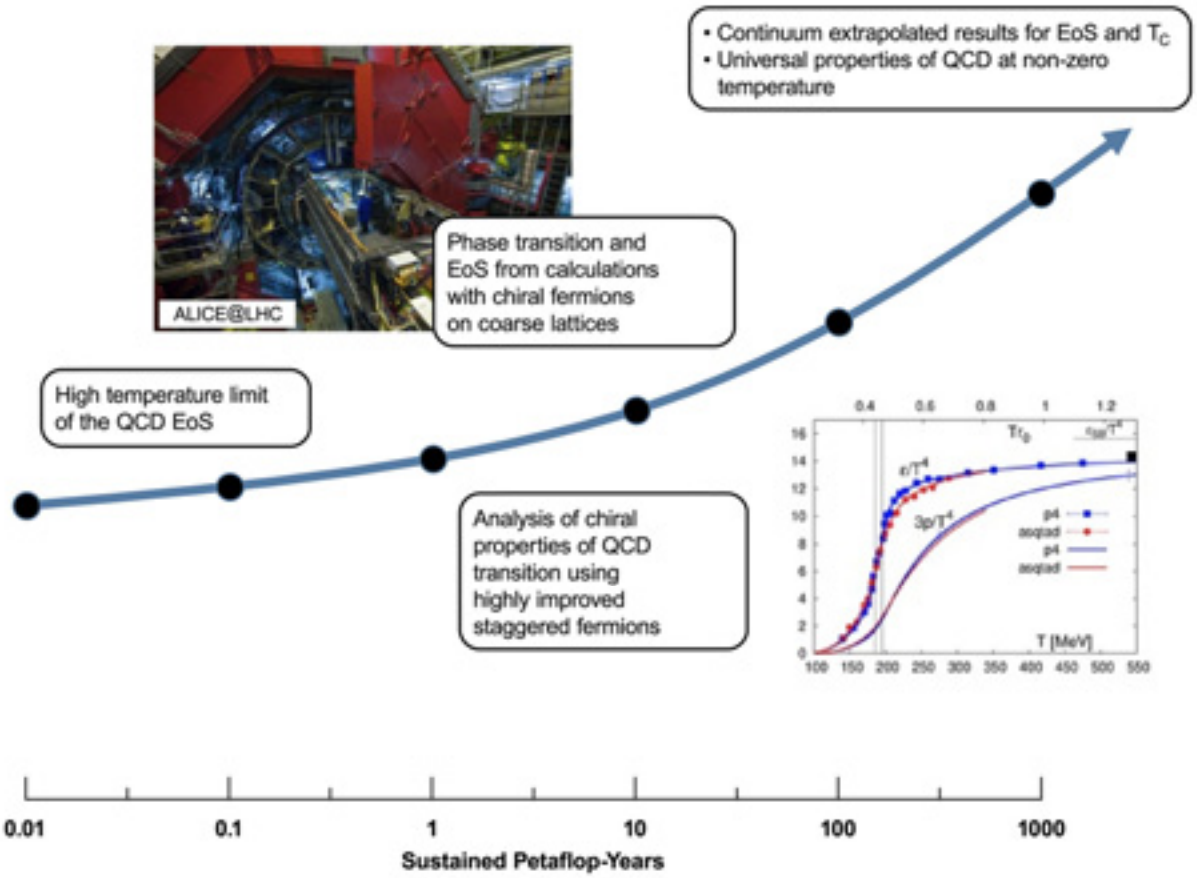

Figure 20. Anticipated highlights for priority research direction "Precision Calculation of Bulk Thermodynamics." Upper-left image courtesy of CERN. Remainder of image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

Scientific Outcomes and Impacts

Establishing the properties of strongly interacting matter at vanishing net baryon number density in the chiral limit will define the anchor point for all studies of the QCD phase diagram as a function of temperature and net baryon number density. It will establish a reliable starting point for extensions of these calculations into the regime of nonvanishing net baryon number density. In combination with calculations using values of light and heavy quark masses as realized in nature, this will quantify the role of chiral symmetry breaking and confinement in the thermodynamics of strongly interacting matter. The equation of state will be the basic equilibrium input to a microscopic description of the rapidly expanding and cooling dense matter formed in a heavy ion collision.

The calculation of the equation of state with physical values of the light-quark masses will not only have a significant impact on the modeling of heavy ion collisions, but it will also constrain the range of validity of conventional perturbative calculations at high temperatures and of model building, based on effective theories, at low temperatures.

Quantum Chromodynamics Phase Structure at Nonzero Net Baryon Number Density

Basic Scientific and Computational Challenges

Current studies of the QCD phase diagram and the thermodynamics at nonzero net baryon number density are limited to the region of small chemical potential; i.e., small net baryon number density. Sensitivity to possible phase transitions at larger values of the chemical potential could arise from conceptually new approaches to the LQCD calculations that overcome the sign problem. This might be achieved through the introduction of auxiliary degrees of freedom that eliminate the oscillating integrands in the QCD partition functions. The complex Langevin approach (Karsch and Wyld 1985; Aarts and Stamatescu 2008) may eventually lead to such an algorithm that avoids the sign problem. However, it has not yet been successfully implemented in realistic calculations. In the absence of such innovative concepts, currently explored techniques will need to be refined to perform calculations with substantially higher numerical accuracy. These numerical approaches include the Taylor expansion of thermodynamic quantities, such as the pressure and energy density, the analytic continuation of results from numerical calculations performed at imaginary baryon chemical potential, as well as approaches that allow for a projection onto physical states with a fixed baryon number. To use these methods in numerical calculations with physical parameters and improved discretization schemes is challenging and goes beyond currently performed exploratory studies.

Taylor Expansion Techniques

To extract sufficient information on the existence of phase transitions in the QCD phase diagram from a series expansion of the QCD partition function (Gavai and Gupta 2003; Allton et al. 2003*)*, which directly gives the expansion of the pressure as function of the baryon chemical potential, many expansion coefficients must be determined. This allows for a systematic analysis of the convergence properties of the series and provides insight into the analytic structure of the partition function. The required numerical effort grows rapidly with the order of expansion. Approximately two orders of magnitude increase in computing resources is required to calculate each additional nonvanishing order in the series expansion.

Analytic Continuation

A straightforward way to avoid the sign problem in calculations at nonvanishing net baryon number density is to replace the baryon chemical potential with a purely imaginary chemical potential (de Forcrand and Philipsen 2002; D'Elia and Lombardo 2003). This enables the use of the highly optimized algorithms developed for the calculation of the QCD equation of state at vanishing chemical potential. In particular, it is possible to perform calculations on large lattices with improved actions. However, to extract information on the thermodynamics at nonvanishing net baryon number density, extremely precise information is needed on the dependence of thermodynamic observables on the imaginary chemical potential. Only then is it possible to analytically continue (i.e. extrapolate) the numerical results to the physically relevant finite density regime.

Canonical Ensemble

An attractive, but extremely computationally demanding approach in the numerical studies of strongly interacting matter at nonzero net baryon number density, is to perform the calculations directly at a fixed

value of the net baryon number density (Kratochvila and de Forcrand 2005; Alexandru et al. 2005). This is in contrast to the approaches discussed above, where calculations are done with an auxiliary control parameter (the baryon chemical potential). To perform calculations in the so-called canonical ensemble generally requires the exact calculation of the determinants of large-sparse matrices. This is straightforward but computationally demanding. Such calculations may profit from improved eigenvalue solvers optimized for QCD applications.

Color Superconducting Phases

At low temperatures, but with large net baryon number density, QCD is predicted to become a color superconductor (Rajagopal and Wilczek 2000; Alford et al. 2008). There may exist several distinct phases, with competing patterns of light-quark flavor-color-spin-momentum pairings. The existence of such phases may have consequences for understanding the evolution of the early universe and the formation of compact stellar objects.

Very little is known from numerical calculations about the phase structure of strongly interacting matter in this regime (away from the extreme asymptotic limits). First-principles calculations in this regime are presently performed only in QCD-like models (Hands 2007). A direct study within QCD will require the development of new techniques that can manage or circumvent the sign problem. Extreme scale computing resources are required to explore such phases.

Computational Challenge

At present, calculations of Taylor expansions up to the third order in the squared baryon chemical potential require about 100 teraflop-years. Extending these expansions to the fifth order will require resources of 1 exaflop-year. To pursue calculations at these high orders, it is necessary to improve the numerical techniques used to calculate Taylor expansion coefficients. Improved techniques for the inversion of large, sparse matrices (deflation) and the optimization of random source vectors (dilution) are currently being tested and are expected to significantly expedite these calculations. The computational challenges that must be addressed in calculations with imaginary chemical potentials are similar. Quantitative studies of finite density QCD, and a decisive calculation that verifies or excludes the existence of a critical point in the QCD phase diagram, require extreme scale computing resources as shown in Figure 21.

Scientific Outcomes and Impacts

Calculations at nonvanishing net baryon number density will greatly advance current knowledge of the phase diagram of strongly interacting matter. High-precision calculations of high-order Taylor expansions, as well as accurate calculations with imaginary chemical potential, will provide information on the analytic structure of the QCD partition function. This may allow definitive statements about the density and temperature dependence of the thermodynamics of dense matter to be made, and eventually may determine the location (or may rule out its existence) of a critical point in the QCD phase diagram.

These calculations will have an enormous impact on current understanding of properties of strongly interacting matter. Further, they will provide strong constraints on the development of theoretical models for the high-density regime of strongly interacting matter, and will influence the accelerator-based experimental research program in this area.

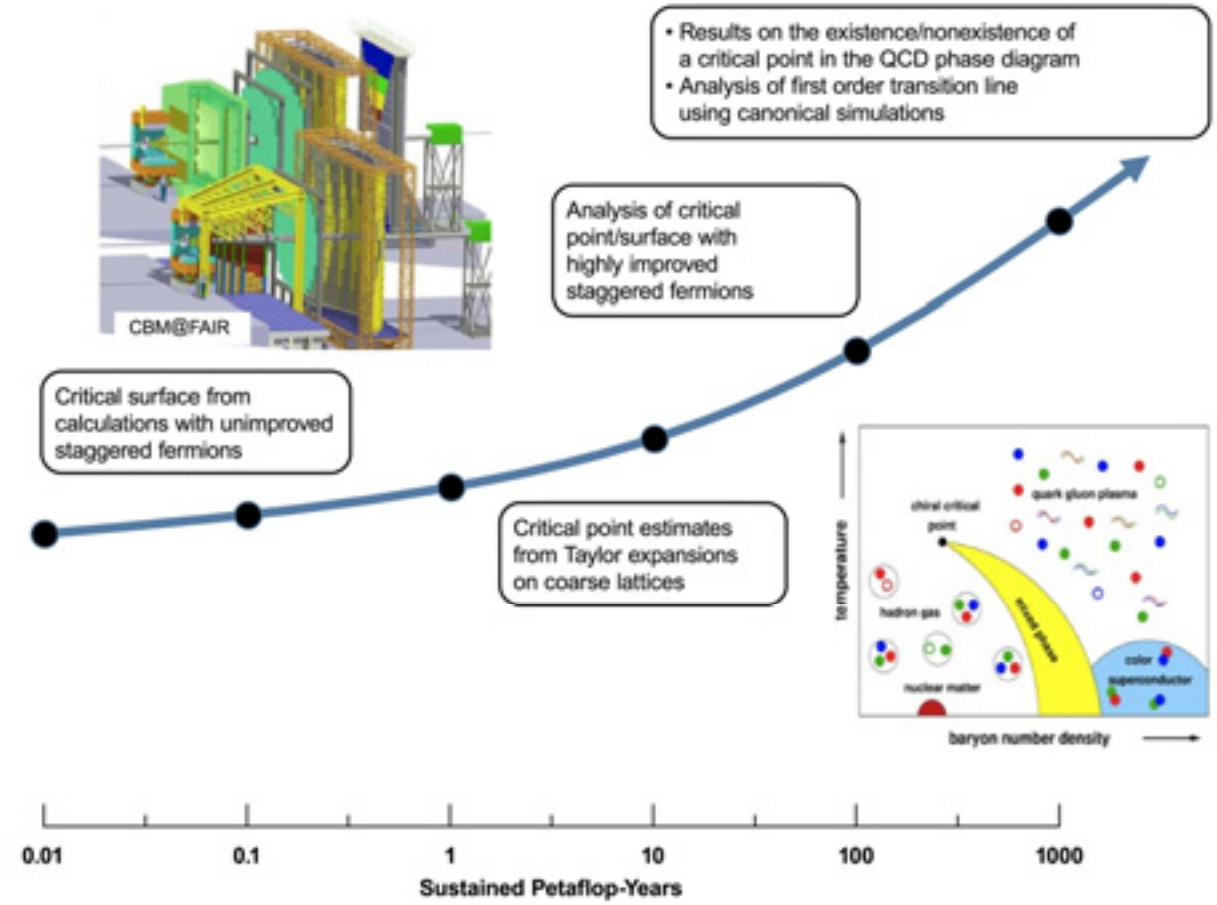

Figure 21. Anticipated highlights for priority research direction "Quantum Chromodynamics Phase Structure at Nonzero Baryon Density." Image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

Transport Coefficients of Quantum Chromodynamics and Spectral Functions of Hadrons in Medium

Basic Science Challenges and Computational Challenges

Numerical calculations of the dynamic properties; i.e., the spectrum of excitations in hot and dense, strongly interacting matter, as well as transport properties of the medium, are presently performed at an exploratory level. To go beyond qualitative statements and to reach a point where quantitative predictions of dynamic properties become feasible, calculations on thermal lattices with unusually large spatial volumes, with greater than 10^3 times the number of lattice sites used in present day calculations, must be performed.

Transport Coefficients

The calculation of transport coefficients, such as the shear and bulk viscosity, that characterize the response of the medium-to-small deviations from its equilibrium state, are particularly difficult. Their calculations formally require taking the limit of zero frequency in an infinite spatial volume, which of

course, is not possible in numerical calculations. To obtain information on the excitation spectrum of a thermal medium requires accurate calculations of correlation functions at a large set of time separations. The extraction of the (continuous) spectral function from a finite set of data points is ill-posed (Karsch and Wyld 1987). To constrain the class of spectral functions that is consistent with these data, the noise level of the data set at the largest time separations has to be below the percent level. Furthermore, the correlation functions have to be calculated at a large number of time separations between sources, making use of correlations between different members of the data set. Thus far, such calculations have only been pursued in quenched QCD (Nakamura and Sakai 2005; Meyer 2007). Even in the quenched case, the lattices that were used were too small to obtain reliable results. A petaflop-year of computing resources will be required to complete the studies of transport properties in quenched QCD, and a computation with light dynamical quark degrees of freedom requires extreme scale computing.

In-Medium Hadron Masses

The degree of difficulty in calculating the hadronic excitations of the medium that provide information on the in-medium modification of light and heavy quark bound states, is similar to that of transport coefficients. Hadrons in a thermal bath interact with the medium, and these interactions can lead to the destruction of bound states, and thus the disappearance of the corresponding resonance peaks in the spectral function. Such an effect has been advocated as an experimental signature for the formation of a hot and dense medium in heavy ion collisions (Matsui and Satz 1986). Indeed, LQCD calculations of spectral functions at high temperature clearly demonstrate the disappearance of resonance peaks from the hadronic spectral functions (Nakahara et al. 1999). However, to follow the disappearance of these states in hot and dense matter in detail, and locate the melting temperature for various hadronic excitations, requires considerably more computing resources than are currently available. Prior to the disappearance of a state, interactions with the thermal medium will lead to temperature and density dependent shifts of the resonance peaks, as well as a broadening of these peaks. To resolve the structure of spectral functions to such a degree that shifts in resonance peaks and broadening of the spectral curve become statistically significant, accurate numerical results for hadron correlation functions are required. As in the case of calculations of transport coefficients, large lattices are needed to generate information on the correlation functions at many different time separations.

Computational Challenge

The major computational challenge in studies of the excitation spectrum of hot and dense matter is the quest for statistically accurate data on correlation functions on large lattices. These lattices are typically a factor of 50 larger than those used in calculations of static, bulk thermodynamics. The size of data samples needed to reach sufficiently small uncertainties in the correlation functions is approximately an order of magnitude larger. Fortunately, such calculations would only be performed at a few selected values of the temperature rather than at the large set of temperature values needed to control properties of the equation of state. Still, this presents a computational challenge, and requires a few petaflop-years to perform calculations within the quenched approximation to QCD. A fully dynamical LQCD calculation, which includes the light-quark contributions, will require extreme scale computing resources, as illustrated in Figure 22.

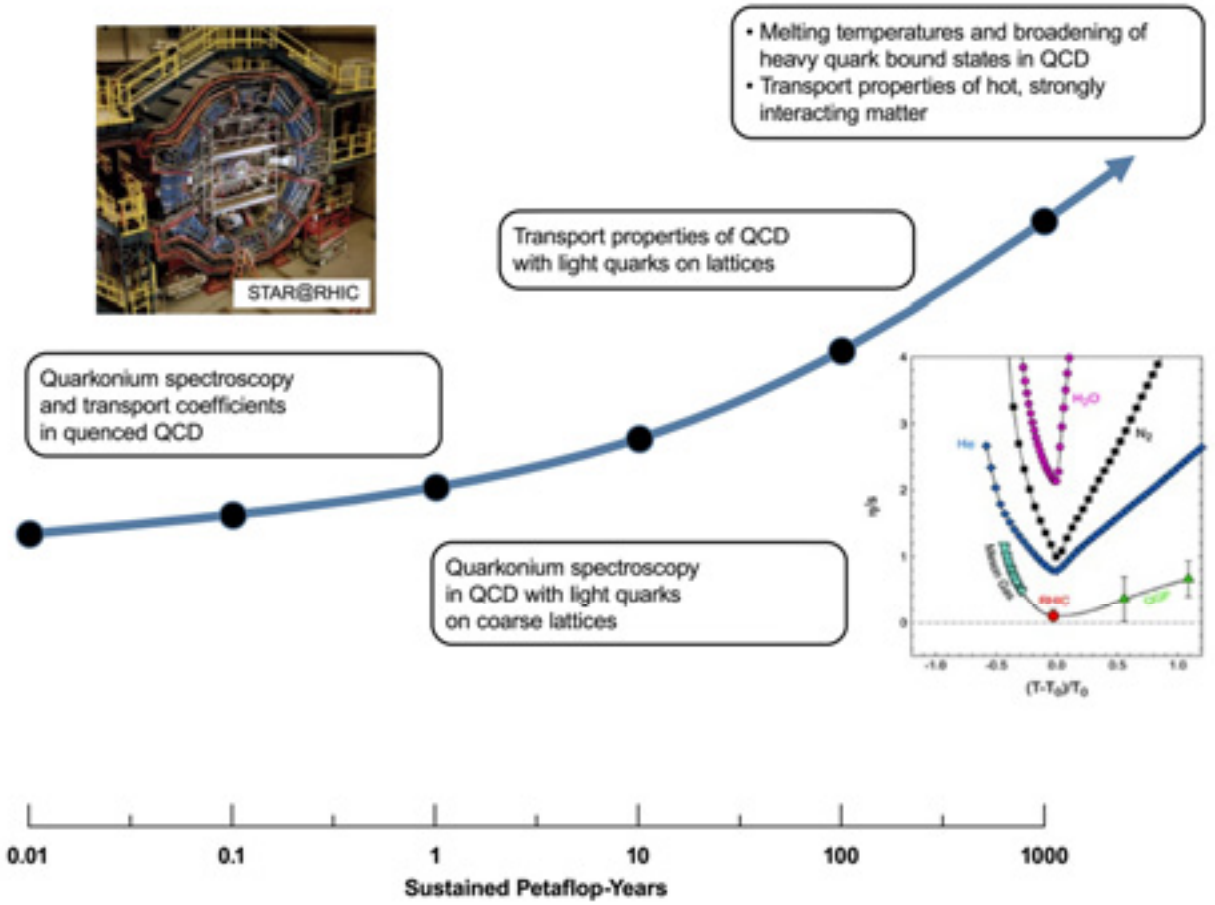

Figure 22. Anticipated highlights for priority research direction "Transport Coefficients of Quantum Chromodynamics and Spectral Functions of Hadrons in Medium." The point labeled RHIC in the right graphic is a theoretical "estimate." Image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

Scientific Outcomes and Impacts

Calculations of transport coefficients will provide fundamental insight into the structure of hot and dense matter. It will allow us to quantify aspects of the extent to which the phenomenologically successful modeling of heavy ion collisions has a solid foundation in QCD; i.e., whether a near-equilibrium QGP described by QCD indeed equilibrates rapidly and can be characterized as an almost-perfect fluid. Detailed information on the spectral function would confirm whether or not the QGP is strongly coupled at RHIC, and by varying the temperature in the LQCD calculations, scientists may learn how much the temperature has to be increased before the plasma becomes weakly coupled. This question will be of importance in comparing the heavy ion data obtained at the RHIC and the LHC experiments because the temperature in the latter will be about a factor 1.5 to 2 higher than in the former.

These calculations will strongly influence the analysis of experimental data obtained in heavy ion collisions.

Equilibration Challenge: From the Color Glass Condensate to the Quark-Gluon Plasma

Basic Science Challenges and Computational Challenges

Determining the time required for a QGP to form after the onset of the collision (i.e., the "thermalization" time) and determining the physics processes that drive the QGP formation are among the most important outstanding problems in the area of relativistic heavy ion collisions. The success of near-ideal hydrodynamics in describing bulk observables—such as the elliptic flow of matter created in noncentral collisions—implies that the matter has a short thermalization time compared to the overall timescale of the reaction.

To describe the approach to equilibrium, the following actions are required: 1) a firm understanding of the initial configuration of partons in the colliding nuclei and the process by which they are liberated from the nucleus at the onset of the collision needs to be acquired; and 2) detailed models and simulations of the processes that occur during the early nonequilibrium phase of the collision leading to the formation of a nearly thermalized QGP need to be developed.

Initial State of the Collision: Color-Glass-Condensate

A large nucleus moving near the speed of light contains a very dense system of gluons. It is believed that nonlinear effects in QCD lead to a saturation of the rapid growth of the gluon density in the colliding nuclei with beam energy and mass number A when the phase-space occupation number is (nonperturbatively) large, on the order of $1/\alpha_S$. The effective theory describing this nonlinear regime of QCD is the color-glass condensate (CGC) (McLerran and Venugopalan 1994a, 1994b; Kovchegov 1996). McLerran and Venugopalan (1994a) proposed an effective action incorporating high-gluon density effects, which amounts to solving the classical Yang-Mills equations where the large-momentum degrees of freedom in the nucleus act as sources of color charge for the small-momentum degrees of freedom.

The classical description of gluon saturation is modified at higher energies due to quantum loop corrections. To this end, a new set of equations, commonly referred to as the JIMWLK (Jalilian-Marian – Iancu – McLerran – Weigert – Leonidov – Kovner) equations (Jalilian-Marian et al. 1997a, 1997b; Iancu 2001), have been derived from a Wilsonian Renormalization Group formalism. They describe the evolution of n-point functions in QCD with energy. The resulting equations are an infinite hierarchy of coupled differential equations that are difficult to solve analytically. Nevertheless, they can be written in a form which, in principle, allows them to be solved by lattice gauge-theory techniques (Rummukainen and Weigert 2004).

The CGC has explained several theoretical and phenomenological aspects of high-energy interactions quite successfully. Nevertheless, many important properties of the CGC remain to be addressed quantitatively and tested by comparison with experimental data from RHIC and other colliders.

Thermalization Mechanisms: Plasma Turbulence

In recent years, it has been shown that early studies of the driving mechanism for the equilibration of quark-gluon matter in ultra-relativistic heavy ion collisions overlooked a crucial aspect of the dynamics of nonequilibrium plasmas—namely, the possibility of plasma instabilities. Most importantly, it has been shown that these instabilities may produce plasma isotropization and approximate thermalization on time-scales relevant to relativistic heavy ion collisions. The possibility of such non-Abelian plasma

instabilities was first predicted in the mid-1990s (Mrówczynski 1993) by studying plasmas with an anisotropic momentum-space distribution. The resulting instability has been dubbed the chromo-Weibel instability. In recent years, this theory has received a significant amount of attention because of analytic and numerical advances. The first advance was to show the instabilities predicted by Mrówczynski are generic and independent of the precise details of the assumed anisotropic parton distribution function (Arnold et al. 2003, 2004; Romatschke and Strickland 2003). Various schemes for treating the non-Abelian, nonequilibrium dynamics (Mrówczynski et al. 2004) using real-time lattice gauge techniques (Hu and Müller 1997) are being explored.

Basic Scientific and Computational Challenges

The full description of the initial state of a heavy ion collision and subsequent thermalization of the matter requires code that can self-consistently describe both the earliest periods when the physics of saturation is important, and also the intermediate times when the physics of the chromo-Weibel instability becomes important. To do this requires the real-time solution of the Yang-Mills equation on three-dimensional lattices coupled self-consistently to the Wong equations. Such codes already exist for simplified configurations and expansion scenarios (Bass et al. 1999; Dumitru et al. 2007; Schenke et al. 2008). The solution of the full problem requires three-dimensional lattices with a fine lattice-spacing in the longitudinal direction; i.e., solving the classical Yang-Mills equations in real time on a three-dimensional lattice with approximately 512^3 sites. The field equations describing the low-momentum gluons need to be coupled self-consistently to the Wong equations that describe the propagation of the hard valence sources in the soft background, including energy-momentum conservation. Beyond the classical limit, a simultaneous solution of the rapidity dependence of the JIMWLK measure together with the real-time evolution of the initial fields is required. Beyond that, it is necessary to have a lattice that is capable of describing the dynamics of the chromo-Weibel instability and the subsequent non-Abelian cascade to high-momentum modes. The "brute force" method to accomplish this is to ensure the lattice spacing is sufficiently fine. This means using lattices significantly larger than 512^3. As some of the processes (e.g., initial conditions, binary particle collisions, and hard radiation) are stochastic, it will be necessary to average observables over multiple sets of initial conditions (runs). Currently, each run of the simplified calculation requires approximately one teraflop-year. Factoring in the higher dimensionality required for the full problem increases this estimate into the tens of petaflop-years region. Averaging over initial conditions and varying experimental parameters will only be possible with extreme scale computing resources as illustrated in Figure 23.

Scientific Outcomes and Impacts

Extreme scale computing will deliver real-time calculations of the collision of two heavy ions at high energy, retaining the complete three-dimensional structure of the fields of produced gluons (in impact parameter and rapidity space), energy-momentum conservation, and quantum evolution of the measure. The subsequent real-time evolution of the color fields following the initial impact will clarify the timescales and processes that lead to thermalization and formation of a QGP and the possible role played by non-Abelian gauge-field instabilities (plasma-turbulence). The distribution of the thermalized gluons in the impact parameter and rapidity will provide much-needed initial conditions for hydrodynamic modeling of the late stages of the collision, and could provide information on the equation of state and the viscosity of hot QCD matter. This work will also provide predictions for the effect of early-time nonequilibrium dynamics on important QGP observables such as jet quenching, anomalous transport, and fluctuations.

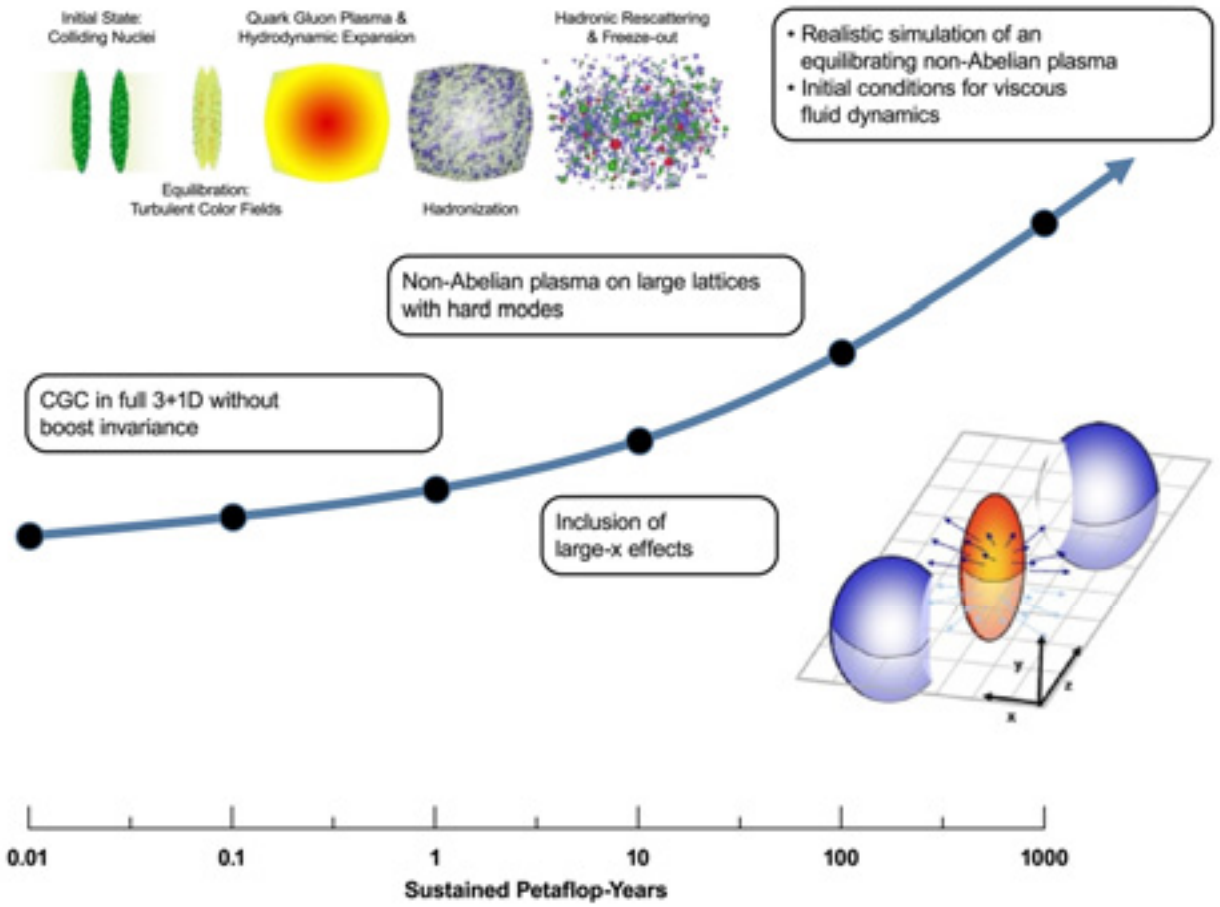

Figure 23. Anticipated highlights for the priority research direction "Equilibration Challenge: From the Color Glass Condensate to the Quark-Gluon Plasma." Lower-right image courtesy of Jerome Lauret. Remainder of image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

APPLIED MATHEMATICS AND COMPUTATIONAL ISSUES

Basic problems, which include numerical calculations in QCD at high temperature and density, are common to other fields that make use of the lattice discretized version of the theory of strongly interacting matter. Current programs for these numerical calculations are quite complex, containing several thousand lines of code. These programs are organized in standardized libraries that have been established during the past few years with the help of support through the Scientific Discovery through Advanced Computing software development initiative.

The central, computationally most demanding part in these calculations is related to the inversion of large, sparse matrices, specifically the fermion matrix (see discussion in the "Cold Quantum Chromodynamics and Nuclear Forces" panel report). Unlike the case in other fields, the specific feature of these sparse matrices, which have on the order of a million rows and columns, is that their nonzero entries fluctuate significantly during the calculation of new field configurations. A specific feature of calculations performed in the analysis of dense QCD is this fluctuating background remains frozen for many inversions and only the source vector used to start the inversion process (e.g., based on a conjugate gradient or a similar algorithm) is varying. The need to invert these matrices several hundred times

underscores the urgency for improved inversion algorithms that can make efficient use of information collected in the previous inversions. Deflation techniques, domain decomposition, and multigrid methods are being tested and implemented for this purpose. These techniques need to be further developed, and in particular, require optimization for new computing architectures that will be used for extreme scale computing applications.

ENHANCED SYNERGISM WITH OTHER SUBFIELDS OF NUCLEAR PHYSICS AND THE BROADER PHYSICS COMMUNITY

Establishing reliable quantitative answers to the physics questions addressed in studies of strongly interacting hot and dense matter is important for the understanding of heavy ion collisions and the basic thermodynamic properties of matter. These calculations are fundamentally important for other subfields of nuclear matter, particle physics, astrophysics, and cosmology. Establishing the structure of the QCD phase diagram and the equation of state of strongly interacting matter will have direct consequences for the modeling of the expansion of the early universe. Future large-scale astrophysics experiments that aim at a detection of primordial gravitational waves may be sensitive to a direct observation of the phase transition of strongly interacting matter. A possible first-order phase transition at high-net baryon number density will influence the modeling of compact stars, and may also lead to observable consequences in the cooling pattern of these stars.

Studies of the phase structure of QCD that are necessary to understand hot and dense matter have many features in common with studies performed in statistical physics and material science. In the past, this led to an engaging exchange of ideas on the level of algorithm development, as well as the development of observables and statistical analysis tools needed for the numerical study of phase transitions in general. Numerical algorithms now used in calculations of strongly interacting matter are based on the Monte Carlo algorithm first applied by Metropolis et al. (1953) in studies of the equation of state of molecules. Current versions rely on the molecular dynamics algorithm first used in chemistry (Anderson 1980) for a similar purpose. Statistical analysis tools like the Ferrenberg-Swendsen algorithm (Ferrenberg and Swendsen 1988) and observables like the Binder-cumulant (Binder 1981) were first applied in statistical physics and are now used as powerful tools to detect phase transitions in strongly interacting matter.

The notorious sign problem faced in numerical studies of QCD at high density led to the development of several new numerical algorithms that have been successfully applied for the simulation of models that are of relevance in statistical physics and material sciences (Chandrasekharan and Wiese 1999).

Numerical calculations of strongly interacting particles have been performed in the framework of LQCD for more than 30 years. These calculations have always demanded large computational resources which, in turn, led several scientists to the development of specific computing hardware. Others developed close collaborations with computer manufacturers. This led to an engaging exchange of ideas and the awareness of the computational needs in LQCD calculations during the design phase of new generations of computers. As a consequence of this close involvement with leading-edge hardware and software developments, many bright scientists who were trained through participation in LQCD calculations now work in the computer and software industry.

CONCLUSIONS

Large-scale computations within the framework of LQCD have played, and will continue to play, a pivotal role in the exploration of the different phases of hot and dense, strongly interacting matter. For many years these calculations exploited leading-edge computing resources that led to a steady improvement of algorithms, as well as discretization schemes used for these calculations. Given this long-standing experience in large-scale computing, the basic tools exist to address the more complex problems that require extreme scale computing resources. Nonetheless, the exploitation of computing resources, which are three to four orders of magnitude larger than those used today for the calculations of hot and dense matter with LQCD, will require the further development and optimization of existing software tools. An increase in computing power by orders of magnitude also opens up possibilities for the exploitation of new discretization schemes—chiral fermions—which have not yet been used efficiently in studies of hot and dense matter. Furthermore, it will elevate calculations of dynamic properties of strongly interacting matter, which can currently only be performed on an exploratory level, to a stage where calculations with physical parameters will provide reliable input to the microscopic modeling of heavy ion collisions and interpretation of experimental results.

Application of real-time lattice techniques to the nonequilibrium evolution of turbulent fields for the dynamic description of the time evolution of a heavy ion collision will require extreme scale computing resources. This fairly new field of super-computing applications will require significant development of software tools to make optimal use of large-scale computing resources.

LQCD calculations are already performed today in large collaborations. The U.S. LQCD community is well organized—it manages its own computing resources, coordinates joint proposals for software development projects, and provides access to leadership class computing facilities. This valuable infrastructure needs to be maintained and extended. Handling computing resources that are orders of magnitude larger than those used today will require increased workloads and a high degree of organization to process the large amount of data generated in these computational projects. Additional new collaborations need to be formed in the area of dynamic modeling to optimize access to the leadership class computing facilities.

The need for access to extreme scale computing resources for a reliable, quantitative study of properties of hot and dense matter is unquestionable.

Table 4 provides an outline of the milestones for the work described in this section. Provided that the computational resources become available for research in hot and dense QCD at the anticipated scales, the forefront research that will be conducted are provided as milestones.

Table 4. Milestones for Hot and Dense Quantum Chromodynamics

Scale	Milestone
>1 Petaflop-year	• LQCD calculations — high-temperature limit of QCD equation of state — critical surface as a function of baryon chemical potential, temperature and light-quark masses using staggered quarks • Color-glass-condensate in 3+1 dimensions with full Lorentz boost-invariance
>20 Petaflop-years	• LQCD calculations — QCD transition and equation of state with chiral quarks on lattices with coarse lattice spacings — quarkonium spectroscopy and transport coefficients in quenched QCD at finite temperature — critical surface as a function of baryon chemical potential, temperature and light-quark masses using highly improved staggered quarks
>100 Petaflop-years	• LQCD calculations at zero net baryon chemical potential — chiral properties of the QCD transition with staggered quarks — quarkonium spectroscopy and transport coefficients with dynamical light quarks at finite temperature • LQCD calculations at nonzero net baryon chemical potential that provide estimates of the QCD critical point on lattices with coarse lattice spacings • Inclusion of large-x effects in the evolution of a heavy ion collisions from a color-glass-condensate to a QGP
>1 Exaflop-year	• LQCD calculations with physical light-quark masses at zero net baryon chemical potential of the following: — continuum extrapolated equation of state and the critical temperature — universal properties of QCD at nonzero temperature — melting temperatures and broadening of heavy-quark bound-states — transport properties of hot, strongly interacting matter • LQCD calculations at nonzero net baryon chemical potential — determination of the existence or nonexistence of the QCD critical-point — analysis of the (possible) first-order transition line using canonical simulations — realistic simulations of an equilibrating non-Abelian plasma • Initial conditions for viscous fluid hydrodynamics

ACCELERATOR PHYSICS

Lead: Robert Ryne, Lawrence Berkeley National Laboratory

INTRODUCTION AND CURRENT STATUS

Accelerators and the U.S. Department of Energy

Particle accelerators are among the most complex and versatile instruments of scientific exploration. Accelerators have enabled remarkable scientific discoveries and important technological advances that span all programs within the U.S. Department of Energy's Office of Science (DOE SC). The importance of accelerators to the DOE SC mission is evident in the DOE report, *Facilities for the Future of Science: A Twenty-Year Outlook* (DOE 2003). Of the 28 facilities listed in that report, half involve accelerators.

The scientific community widely recognizes that particle accelerators are essential tools for the study of elementary particles. Among the list of Nobel Prizes in physics since 1939, more than a dozen of the prizes awarded involved discoveries at particle accelerators or the development of accelerator and detector technologies. It is useful to review the accelerator facilities operated by the DOE SC, and to describe the extreme importance of particle accelerators to U.S. science and technology (see sidebar).

Along with advances in nuclear and particle physics, modern accelerators are also crucial to advances in materials science, chemistry, the biosciences, and other fields. Many of these advances are made at the nation's spallation neutron sources and synchrotron light sources. These applications are incredibly wide ranging and for example include determining the structure of proteins; imaging and determining the structure of organelles in cells; exploring high temperature superconductivity; designing plastics and most modern composite materials; determining the structures of magnetic materials now ubiquitous in computer and telecommunications technology, and designing medicinal drugs.

In addition to applications in basic and applied science, accelerators have applications in national security, energy and environmental security, health, and medicine. Regarding national security, accelerators are used for stockpile stewardship; e.g., in applications involving neutron and proton radiography. Applications to the Department of Homeland Security's mission include compact, accelerator-based neutron generators for screening sea-bound cargo containers, airport containers, and vehicles at border crossings. Active interrogation systems are under development to detect special nuclear materials, explosives, and other contraband.

In energy and environmental applications, accelerator-driven fission energy production systems have been proposed that have the potential to be both safe and environmentally friendly. Such systems would use a subcritical assembly (as opposed to a critical assembly as found in nuclear reactors) that could potentially burn a large fraction of its own byproducts, thereby mitigating waste and waste disposal issues. Systems have also been proposed for the accelerator transmutation of waste in which waste containing very long-lived radioisotopes is transmuted to shorter-lived byproducts that are much easier to store, and need to be stored for a much shorter period of time. On energy production, the United States has made great strides in accelerator-based heavy ion fusion under research and development programs supported by the DOE Office of Fusion Energy; this research has important applications to high-energy density physics research.

Particle Accelerators in Science and Society

The scientific community widely recognizes that particle accelerators are essential tools for the study of nuclei and elementary particles. Among the Nobel Prizes awarded in physics since 1939, more than a dozen prizes involved discoveries at particle accelerators, or the development of accelerator and detector technologies.

Along with advances in nuclear and particle physics, modern accelerators are also crucial to advances in materials science, chemistry, the biosciences, and other fields. Many of these advances are made at the nation's spallation neutron sources and synchrotron light sources. These applications are incredibly wide ranging and can be used in determining the structure of proteins, imaging and determining the structure of organelles in cells, exploring high-temperature superconductivity, and designing medicinal drugs.

In addition to applications in basic and applied science, accelerators have applications to national security, energy and environmental security, healthcare, and medicine. Regarding national security, accelerators are used for stockpile stewardship; e.g., in applications involving neutron and proton radiography. Applications to the Department of Homeland Security's mission include compact, accelerator-based neutron generators for screening sea-bound cargo containers, airport containers, and vehicles at border crossings. In energy and environmental applications, accelerator-driven fission energy production systems have the potential to be safe, environmentally friendly, and proliferation resistant. Systems have also been proposed for the accelerator transmutation of radioactive waste. For energy production, the United States has made great progress in accelerator-based heavy ion fusion under research and development programs supported by the U.S. Department of Energy Office of Fusion Energy; this research also has important applications to high-energy density physics research. In the health and medicine areas, it is estimated that about 10,000 cancer patients are treated daily in the United States with electron beams from linear accelerators. Medical accelerators that accelerate hadrons have also been developed and are especially useful for irradiating deep-seated tumors. Accelerators are also used to produce radioisotopes for medical treatment and diagnosis. This can prove critical where chemotherapy methods cannot target the affected site.

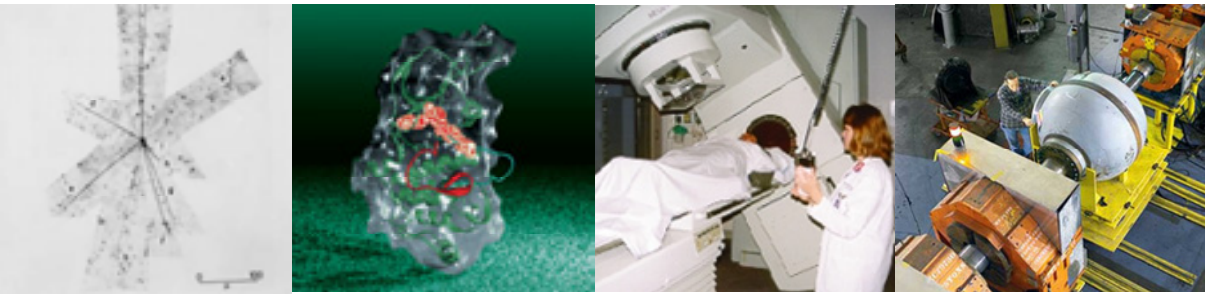

Particle accelerators are among the most versatile and important tools of discovery in basic and applied research. These accelerators have a huge impact on progress in U.S. science and technology, and consequently on the U.S. economy and our quality of life. Images from left to right: discovery of the antiproton; accelerator-based protein crystallography; accelerator-based medical therapy; and proton radiography for stockpile stewardship. Images courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

In health and medicine areas, it is estimated that about 10,000 cancer patients are treated daily in the United States with electron beams from linear accelerators. Medical accelerators that accelerate hadrons have also been developed and are especially useful for irradiating deep-seated tumors. Accelerators are also used to produce radioisotopes for medical treatment and diagnosis.

In summary, particle accelerators are among the most versatile and important tools of discovery in basic and applied research. These instruments have a huge impact on progress in U.S. science and technology, and consequently on the economy and our quality of life.

Accelerators are also objects of research for the scientific community in their own right. Demands for high energy, high intensity, low-loss rates, low activation of components, low cost, ease of tuning, and prediction of performance in new regimes all need considerable theoretical, analytical, and computational work. Novel designs and those employing novel materials and/or materials pressed to their performance limits require significant analysis before the first prototype is built. New technologies involving laser- and plasma-based acceleration have the potential to revolutionize accelerator technology. However, the physical phenomena involved are extremely complex; large-scale computing, in concert with theory and experiment, is essential to understand these phenomena and fully develop these new technologies. Optimization of operation of accelerators also requires considerable analysis.

Advanced computing, particularly the Scientific Discovery Through Advanced Computing (SciDAC) program, has had a major impact on the field of accelerator science and technology. Examples of active areas of research, where advanced computing is essential to develop the physics, engineering, and technologies of future systems include the following: electron cooling and stochastic cooling systems, which preserve and enhance luminosity and lifetime for existing facilities; systems to maximize long-term storage of beams in colliders; the design of heavy ion fragment separators to extract very rare produced isotopes that may be present at only one part in 10^{13} or less in the reaction products produced during bombardment; technologies to reduce detector backgrounds in both colliders and fixed-target machines, by reducing population of phase-space outside of desired areas; and technologies for the production and handling of bremsstrahlung and beamstrahlung in electron machines.

Currently, DOE SC operates about a dozen large accelerator facilities. The major accelerator user facilities of the Office of Nuclear Physics (Figure 24) include the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL), the Continuous Electron Beam Accelerator Facility (CEBAF) at the Thomas Jefferson National Accelerator Facility (Jefferson Laboratory), the Argonne Tandem Linac Accelerator System at Argonne National Laboratory (ANL), and the Holifield Radioactive Ion Beam Facility at Oak Ridge National Laboratory (ORNL).

The main facility of the Office of High Energy Physics (Figure 25) is the Tevatron complex at the Fermi National Accelerator Laboratory.

The Office of Basic Energy Sciences operates five light sources: the Linac Coherent Light Source and the Stanford Synchrotron Radiation Laboratory at the Stanford Linear Accelerator Center National Accelerator Laboratory; the National Synchrotron Light Source at BNL; the Advanced Photon Source at ANL; and the Advanced Light Source at Lawrence Berkeley National Laboratory (LBNL). The Office of Basic Energy Sciences also operates the Spallation Neutron Source at ORNL, and provides support to activities at the Lujan Center at the Los Alamos Neutron Scattering Center facility (see Figure 26).

Figure 24. Aerial photos showing RHIC at Brookhaven National Laboratory (top left); CEBAF at the Thomas Jefferson National Accelerator Facility (top right); the Argonne Tandem Linac Accelerator System at Argonne National Laboratory (lower left); and the Holifield Radioactive Ion Beam Facility at Oak Ridge National Laboratory (lower right). Images courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Figure 25. Aerial photo of Fermi National Accelerator Laboratory. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Figure 26. Top-left photo is the Stanford Linear Accelerator Center National Accelerator Laboratory, home of the Linac Coherent Light Source and Stanford Synchrotron Radiation Laboratory; top-right photo is the Advanced Light Source at LBNL; middle-left photo is the National Synchrotron Light Source at BNL; middle-right photo is the Advanced Photon Source at ANL; and the lower-left photo is the Spallation Neutron Source at ORNL. Images courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Accelerators and the DOE Office of Nuclear Physics

The mission of the Nuclear Physics program is to promote nuclear physics research through the development and support of basic research scientists and facilities. Nuclear physics research seeks to understand the fundamental forces and particles of nature as manifested in nuclear matter. DOE provides about 85% of U.S. nuclear physics research funding, much of which is directed toward the development, construction, and operation of large state-of-the-art accelerator facilities and detectors at national laboratories. These facilities are used by researchers from laboratories and universities supported by DOE, the National Science Foundation, other agencies, and foreign countries.

Nuclear physics accelerator facilities supported by DOE at national laboratories include the RHIC at BNL, the CEBAF at Jefferson Laboratory, the Argonne Tandem Linac Accelerator System at ANL, the Holifield Radioactive Ion Beam Facility at ORNL, and the 88-Inch Cyclotron at LBNL. Facilities at

universities include the Cyclotron Institute at Texas A&M University, the Wright Nuclear Structure Laboratory at Yale University, the Triangle Universities Nuclear Laboratory at Duke University, and the Center for Experimental Nuclear Physics and Astrophysics at the University of Washington. All these facilities have led to important scientific advances and discoveries in nuclear physics. Along with existing facilities, planning for new and proposed facilities is underway. These include the CEBAF 12 GeV upgrade, now underway at Jefferson Laboratory; the Facility for Rare Isotope Beams (FRIB) at Michigan State University; and proposals for electron-ion colliders (EICs)—the EIC concept at RHIC (eRHIC) and the Electron Light Ion Collider (ELIC) concept at Jefferson Laboratory.

The design, optimization, operation, and upgrade of accelerator facilities involve an enormous amount of advanced scientific computing and simulation. Examples include the self-consistent simulation of beam-beam effects, electron-cloud effects, electron cooling channel design, ion source design, and the design of complex three-dimensional electromagnetic structures. Such simulations allow scientists to understand complex beam and electromagnetic phenomena, to test new ideas, optimize designs, and reduce cost and risk. Ultimately, such simulations play a critical role in the success of accelerator projects.

Nuclear Physics Priorities for Future Accelerator Facilities

The Frontiers of Nuclear Science: A Long Range Plan (DOE 2007) makes three recommendations regarding accelerator facilities:

1. We recommend completion of the 12 GeV CEBAF Upgrade at Jefferson Lab. The upgrade will enable new insights into the structure of the nucleon, the transition between the hadronic and quark/gluon descriptions of nuclei, and the nature of confinement.
2. We recommend construction of FRIB, a world-leading facility for the study of nuclear structure, reactions, and astrophysics. Experiments with the new isotopes produced at FRIB will lead to a comprehensive description of nuclei, elucidate the origin of the elements in the cosmos, provide an understanding of matter in the crust of neutron stars, and establish the scientific foundation for innovative applications of nuclear science to society.
3. The experiments at the Relativistic Heavy Ion Collider have discovered a new state of matter at extreme temperature and density—a quark-gluon plasma that exhibits unexpected, almost perfect fluid dynamical behavior. We recommend implementation of the RHIC II luminosity upgrade, together with detector improvements, to determine the properties of this new state of matter. (DOE 2007)

Under the heading, "The Emerging QCD Frontier: The Electron-Ion Collider," the Nuclear Science Advisory Committee report (DOE 2007) presents a scientific justification for an EIC:

> To develop the most compelling case for EIC in a timely way it is clear that over the next five years significant progress must be made in the conceptual design of the accelerator. It will be important to converge on one accelerator concept so that detailed plans and schedules can start to be developed. This can only happen after essential research and development is completed in a number of areas including: cooling of high-energy hadron beams, high-intensity polarized electron sources, and high-energy, high-current Energy Recovery Linacs. This research and development effort should be carried out so as to leverage existing expertise and capabilities at laboratories and universities. Research and development for a high-energy EIC will maintain U.S. leadership in an area with important societal applications. The design of collider detectors integrated into the accelerator will also be a key issue requiring detailed physics simulations as well as detector research and development. (DOE 2007)

Along with these four priorities (CEBAF 12 GeV upgrade, FRIB, RHIC II, and an EIC) the DOE (2007) report strongly affirms the importance of accelerator research and development:

> Discoveries in nuclear science rely significantly on particle accelerators. Progress in accelerator science and technology is essential for the development of new capabilities that will enable future discoveries. (DOE 2007)

Furthermore, the DOE (2007) report highlights two technologies that the Nuclear Physics program helped pioneer: Energy Recovery Linac technology, and Energy Recovery Linac-based Free-Electron Laser (FEL) technology.

Current Capabilities in High-End Simulation of Accelerators and Accelerator Components

The application of parallel computing in accelerator modeling began in earnest in the 1990s. Starting with a DOE Grand Challenge in computational accelerator physics, the scientific community effort grew into the SciDAC-1 project known as "Advanced Computing for 21st Century Accelerator Science and Technology." The current effort is the "Community Petascale Project for Accelerator Science and Technology," or ComPASS, funded under SciDAC-2, the follow-on to SciDAC-1. The ComPASS project has participants from all DOE national laboratories with major accelerator facilities, as well as participants from universities and small businesses. The main focus of ComPASS has been to develop a suite of multiphysics design codes able to model complete accelerator systems and able to scale to tens of thousands of computing cores or more. The codes fall mainly into three categories: beam dynamics, electromagnetics, and advanced (laser/plasma) accelerator design codes.

Compared with the early parallel accelerator codes of the 1990s, ComPASS codes are able to perform simulations roughly 100,000 times more challenging. For example, beam-dynamics codes of the 1990s used tens of thousands of simulation particles with simplified (e.g., two-dimensional) models of collective effects; currently, such simulations can be performed with three-dimensional models, with more than a billion particles, and with more physical phenomena included. Similarly, current simulation codes are able to design complex three-dimensional accelerator structures with accuracy greater than fabrication tolerance. Such electromagnetic modeling codes have dramatically reduced the design cycle time for accelerator structures. Furthermore, the optimization of electromagnetic structures through high-resolution modeling will reduce the overall cost of future facilities.

The following illustrates several examples of previous large-scale parallel simulations for nuclear physics accelerator projects.

Design of Facility for Rare-Isotope Beams

Detailed design optimization and beam-dynamics simulations for the FRIB accelerator systems have been performed by groups from ANL, Michigan State University, LBNL, and Los Alamos National Laboratory using parallel beam dynamics codes (IMPACT/RIAPMTQ and TRACK) running on computers at the National Energy Research Scientific Computing Center, the High-Performance Computing Center at Michigan State University, and at ANL. These simulations included various machine errors to evaluate beam losses, which involved multiple design revisions to obtain a robust and cost-effective design of the facility. Figure 27 shows an example of beam envelopes along the driver linac from an error study that used 100 realizations of the accelerator lattice (i.e., from a study whose errors were generated from 100 random seeds). Each case involved simulating a uranium beam of 2×10^5 particles. Results show that the final optimized design is robust and tolerant to a typical set of machine errors without any uncontrolled beam losses.

Computers and Accelerators in History

The 1930s may be viewed as the birth of particle accelerators, with such inventions as the Cockroft-Walton accelerator; the Van de Graaff generator; the linac, and the cyclotron. The 1940s may be regarded as the birth of computers with the invention of the first fully electronic computers. Over the decades, these technologies have advanced side-by-side. Accelerators progressed from million electron-volt (MeV) energy to billion (GeV) to trillion (TeV), while computers advanced from megaflops to gigaflops to teraflops to petaflops.

Early on, accelerator physicists recognized the importance of computers. Particle tracking calculations were performed using the Illinois Automatic Computer in the 1950s. Over the years, computers became widely used for tracking and for magnet and radio frequency cavity calculations. A prescient comment is found in the 1971 Proceedings of the 8th International Conference on High Energy Accelerators. Following the opening address, there was a question and answer session. One comment was from Lew Kowarski, who played a key role in the founding of the European Organization for Nuclear Research (CERN):

> "I would like to comment on your three kinds of physicists [experimentalists, theorists, machine physicists] in a perspective somewhat more extended in time. Early experimentalists worked with their hands: Galileo's legendary tossing of stones from the Tower of Pisa, or the alchemists mixing by hand the ingredients in their mixing bowls. In a similar way the theoreticians manipulated their numerical quantities and symbols by their unaided brain-power. Then came the machines to extend the experimenter's manual skill and to open whole new worlds of things to be handled in ways nobody could predict or even imagine before they really got going. Now we are at the beginning of a new kind of extension by machine: the computer comes to supplement the theoretician's brain. We cannot foresee what this fourth kind of creativity in physics will bring, but we may expect that, just as Ernest Lawrence's contribution was decisive to the development of nuclear machines, the name of John von Neumann will be remembered in connection with the origins of computational physics."

Images from left: Ernest Lawrence, Lew Kowarski (middle image, second from left), and John von Neumann. Images courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Scientific computation has been essential to the development of modern particle accelerators. Kowarsky's comments about "the beginning of a new kind of extension by machine" were made in 1971 at a time when the fastest computer in the world was the CDC 7600. Codes on the CDC 7600 could achieve performance of around 10 megaflops. Scientists are now using petaflop computers that are 100 million times more powerful than the CDC 7600.

Accelerator physicists and DOE continue to recognize the importance of accelerators to science and society, and the enormous impact of scientific computing on the development of accelerator science and technology. Advanced scientific computing is essential to study collective phenomena and mitigate beam instabilities, to design complex three-dimensional electromagnetic structures, and to explore new methods of acceleration. Ultimately, advanced computing is essential to optimize designs, to reduce cost and risk, and to help assure the success of accelerator projects.

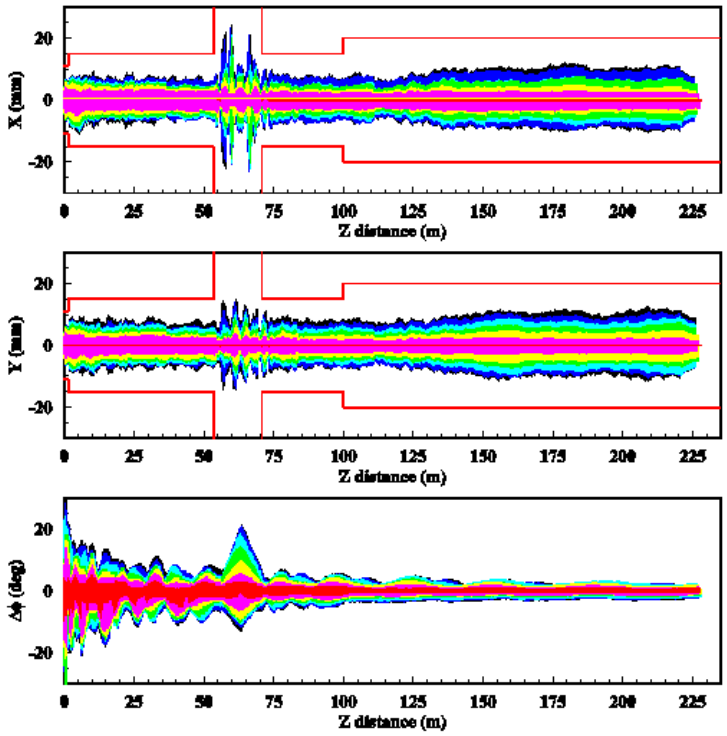

Figure 27. Uranium-beam envelopes along the FRIB driver linac for 100 seeds of machine errors with 2×10^5 particles for each seed. The red lines in the top two plots show the actual aperture of the linac. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Design of a High Resolution Isotope Separator for Facility for Rare Isotope Beams

One of the most important aspects of experimenting with rare isotopes is related to the separation purity of the isotopes to be selected. If the reaction mechanism involves fragmentation and/or fission of heavy ion primary beams, then a multistage isotope separator is employed. The separator must be designed such that the primary beam rejection is perfect, and the several hundreds of other species that act as unwanted impurities for the experimental stations are minimized. The tiny production cross-sections of rare isotopes of interest to nuclear physics make this an especially difficult problem to model and simulate accurately and efficiently. The computational challenges are dependent on the rarity of the isotopes to be studied. These challenges include research currently possible on a typical single-processor personal computer (such as ^{14}Be obtained from fragmentation of oxygen requiring approximately 10-gigaflop days) and cases that are manageable with petaflop machines (such as ^{132}Sn shown in Figure 28 or ^{78}Ni from the fission of uranium requiring approximately one-petaflop day). Future cases for FRIB will need more powerful computers because FRIB is being built to study nuclei with even more extreme neutron-to-proton ratios and therefore are rarer among the fragmentation products. This is discussed further in the priority research direction (PRD) section.

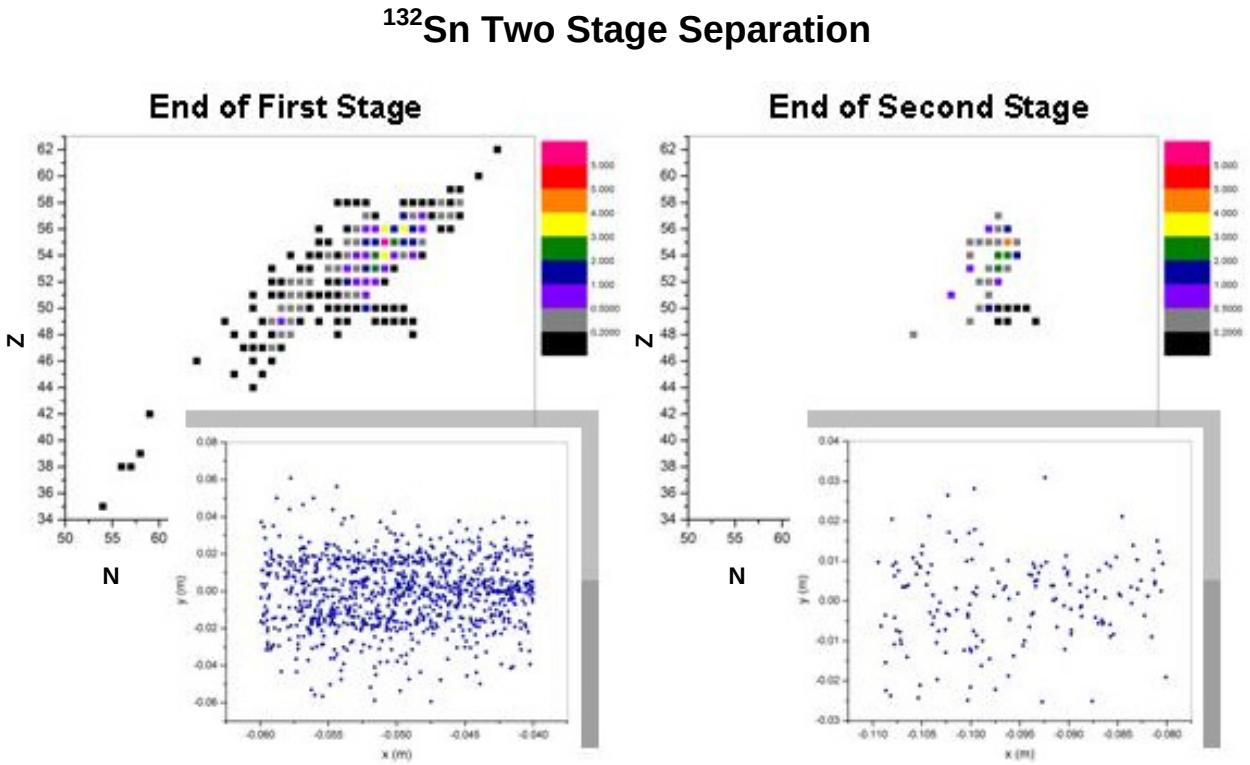

Figure 28. Separation purity of ^{132}Sn obtained from the fission of a 200 MeV/u ^{238}U primary beam using a proposed two-stage isotope separator. The purity plots are illustrated in the N-Z chart of nuclei, and the insets are the expected x-y focal plane distribution of ^{132}Sn ions after one and two separation stages, respectively. The color coding in the purity plots is obtained by normalizing the intensity of the impurities by the intensity of the tin beam. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Design of a Future Electron-Ion Collider

Beam-beam effects are a critical phenomenon affecting the luminosity—and hence the scientific discovery potential—of colliders. As such, beam-beam studies are essential for a successful conceptual design of a future electron-ion collider as well as for operational optimization of an existing collider. Because of the highly nonlinear and collective nature of the problem, large-scale computer simulation is currently the only viable approach to design optimization. Great progress has been achieved in the last two decades in both the physics understanding and the simulation methodology for beam-beam studies, with codes being developed worldwide and used for the design and operation of colliders. Presently, computer simulations are generally carried out using a strong-strong (i.e., self-consistent) model with transverse beam-beam force calculations, the inclusion of multiple physical phenomena, and use of large numbers of simulation particles. Large-scale beam-beam simulations have supported existing facilities (RHIC) and the design of future facilities.

In regard to a future EIC, two concepts have been analyzed: a linac-ring design (denoted eRHIC) being studied at BNL and a ring-ring design (denoted ELIC) being studied at the Jefferson Laboratory (see Figure 29 and Figure 30). Simulation studies have concentrated on incoherent and coherent instabilities, beam-emittance degradation, and luminosity reduction as a function of machine parameters.

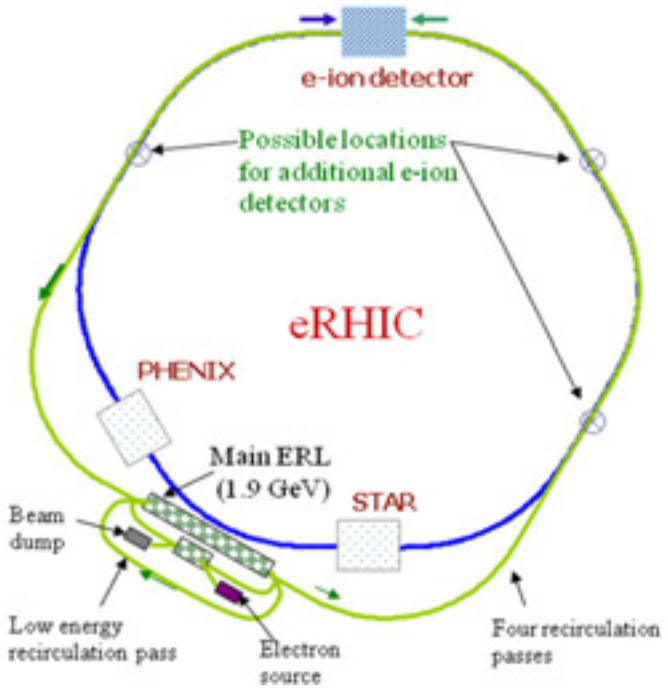

Figure 29. A schematic drawing for the Electron Relativistic Heavy Ion Collider (eRHIC) Linac-Ring Concept. PHENIX and STAR are the present large experiments at RHIC. The energy-recovery linac (ERL) would be a key addition to the present RHIC. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Figure 30. A schematic drawing for the Electron Light Ion Collider (ELIC) ring-ring concept. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

As an example, Figure 31 shows the expected luminosity in ELIC based on an early working point (i.e., an early set of parameters in a design space) and an improved working point studied through large-scale simulation. In this example, large-scale simulation led to a new working point that offers approximately 30% improvement in luminosity and therefore a similar gain in discovery potential.

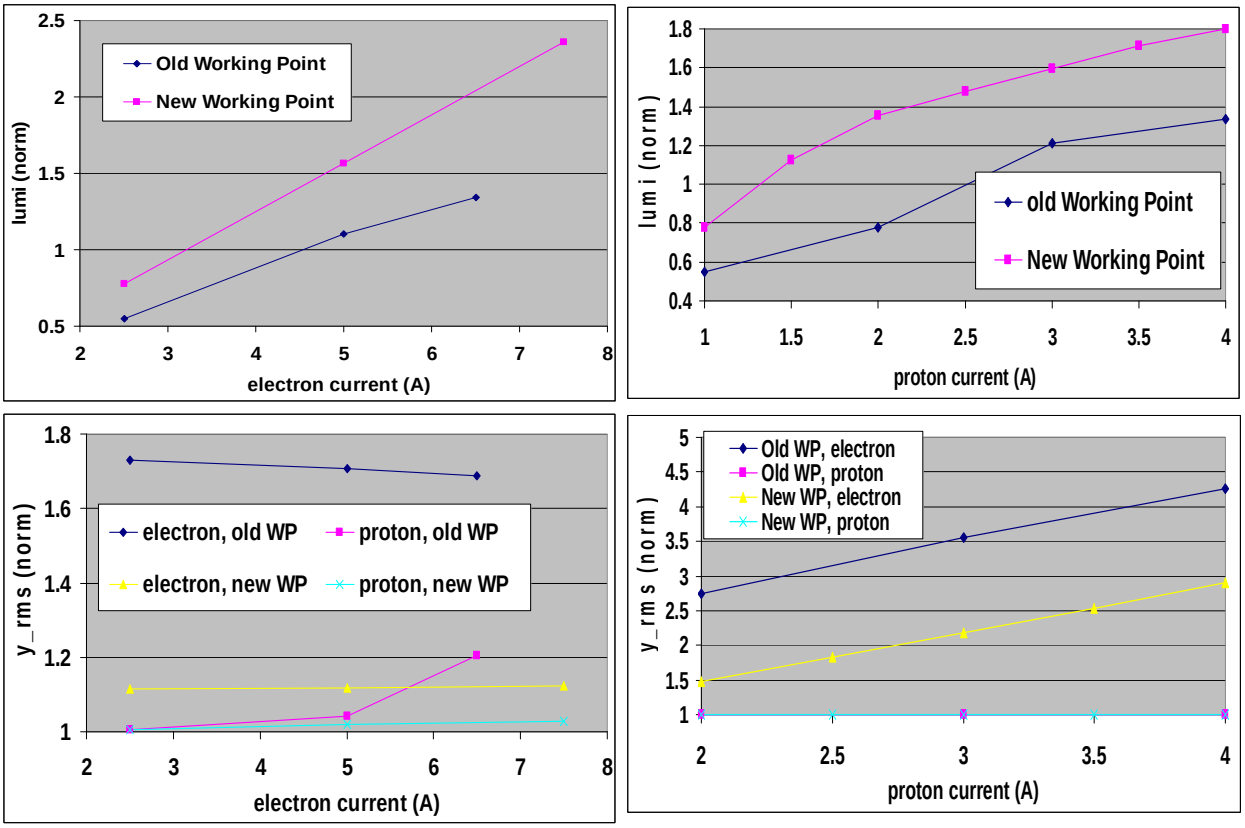

Figure 31. Luminosity as a function of electron (left) and proton (right) current at the ELIC proposed by the Thomas Jefferson National Accelerator Facility and based on simulation using the Beambeam3d Parallel Beam Dynamics code. Over much of the parameter space, the new working point (WP) found through large-scale simulation offers 30% improvement in luminosity and therefore a similar gain in scientific discovery potential.

Simulation of Beam-Beam Effects at the Relativistic Heavy Ion Collider

Multiple bunch self-consistent strong-strong beam-beam simulations of RHIC have been conducted on high-performance computers using the BeamBeam3D code (Qiang et al. 2002). The purpose of the simulations was to find a new working point (i.e., new set of magnet settings) to optimize the beam quality and hence the luminosity. Near the nominal working point and near a half integer in tune space, emittance growth of more than 100% was observed in the parameter space. Parallel simulations with BeamBeam3D were used to study a new working point near an integer resonance. A new working point was found for which the emittance growth was below 10% in much of the parameter space. These large-scale parallel simulations will improve future RHIC operations.

Simulation of the Beam Breakup Instability in the Thomas Jefferson National Accelerator Facility Free-Electron Laser

Recently, scientists performed the first self-consistent beam breakup simulations of an energy recovery linac including multiple passes in a full nine-cell structure for parameters relevant to the Jefferson Laboratory FEL (Smithe 2008). These were the first to couple electromagnetic simulations of the accelerating cavity with particle tracking to model the FEL dynamics and time delays of bunch re-entry. The simulations showed self-consistent generation of higher-order modes by the beam on the first pass

and subsequent interaction of the beam with those modes on the second pass. Between the first and second passes, both an energy decrease and an energy spread of a few percent together modify the beam phase space, to model the effect of a free-electron laser. Analysis of the cavity's resulting mode spectrum showed that the first pass of the beam generated roughly 50 measurable higher-order modes, and scientists identified parameter regimes where modes would grow after multiple passes of the beam. These simulations required resources of the order of a teraflop-hour; even so, scientists were required to use an artificially large simulated transverse beam size and did not include the detailed geometry for mode damping due to limited computing capability. Increased computing resources are required to perform simulations with realistic geometries and beam size.

Design of Electron Cooling Systems for Nuclear Physics Applications

Novel EIC concepts are a high priority for the long-term plans of the international nuclear physics community. Orders-of-magnitude larger beam transverse phase space density than that achieved at the RHIC, or planned at the Large Hadron Collider, is required for the relativistic ion beams in such accelerators. Electron cooling is a promising approach to achieve the necessary luminosity. Many low-energy electron cooling systems have been successfully built and operated for nonrelativistic ion storage rings, and recently a moderately relativistic cooling system was successfully commissioned at Fermi National Accelerator Laboratory (Nagaitsev et al. 2006). Cooling systems have not been built for EIC parameters, but simulations using high-performance computers with the VORPAL®[1] framework (Nieter and Cary 2004) have resolved differences in analytical models for the dynamical friction force in magnetized cooling systems (Fedotov et al. 2006). These (parallel) simulations have also quantitatively shown how the friction would be reduced if the typical high-field solenoid magnet is replaced with a conventional (and much less expensive) undulator (Bell et al. 2008). As shown in Figure 32, the friction force is reduced logarithmically with the undulator field strength.

Stability Proofs of Nonlinear Dynamics in Synchrotrons

A key issue for synchrotrons and colliders is the stability of particle motion inside the beam pipe. Using a one-turn transfer map of sufficiently high order obtained through differential algebraic methods (Berz 1999), it is possible to determine a nonlinear coordinate transformation to a normal form in which the motion is nearly circular. The slight deviations from circularity, the so-called normal form defect, allow a rigorous bounding of the stability times of the nonlinear motion and hence the stability of the underlying synchrotron or other repetitive accelerator. An example of a normal form defect is shown in Figure 33, which illustrates the rich structure of local extrema that may be present. Using rigorous global optimizers, it is possible to determine bounds with an accuracy that is sufficient to establish stability times (Berz et al. 2005). This accuracy would provide a robust method to support the setup and tuning of all large accelerators, because the long-term stability of the orbit is necessary to find a stable operating point, to decrease backgrounds in detectors, and to reduce radioactivation of accelerator components.

Because of the small scale and exceedingly complicated structure of the normal form defect, this stability calculation is a challenging problem usually performed on clusters of several thousand cores—yet the technique has only been applied to existing machines with fixed parameter choices. The logical extension, currently impossible because of limited available computational resources, is to combine the stability estimates with parameter optimization to perform a global search that maximizes stability times. This re-iterates the importance of global optimization tools.

[1] VORPAL is a registered trademark of the University of Colorado.

Impact of Advanced Computing on Accelerator Science and Technology

Particle accelerators and their associated detectors are among the largest, most complex scientific instruments in the world. The successful development of large accelerator facilities involves enormous investments in the three paradigms of scientific research: theory, experiment, and computation. Neglecting any of these can lead to an inability to meet performance requirements, cost overruns, and ultimately project failure. For example, in the early 1990s, uncertainty in the design aperture of the proposed Superconducting Super Collider led to a decision to increase the aperture by 1 cm. This led to a projected $1 billion cost increase to the project, which was eventually cancelled.

In contrast, current advanced computing diagnoses issues before they present serious consequences to projects. An example is provided by the 12 GeV upgrade at Jefferson Laboratory. Beam breakup at well below the designed beam current was observed in the Continuous Electron Beam Accelerator Facility 12 GeV upgrade prototype cavities. Higher order modes (HOMs) with exceptionally high-quality factors were measured as seen in Figure 1(a). The cause was attributed to cavity deformation during the fabrication process. Using the measured cavity parameters as input, the deformed cavity shape was recovered by solving the inverse problem through an optimization method developed by researchers at the Stanford National Accelerator Laboratory. The calculations showed that the cavity was 8 mm shorter than designed, which was subsequently confirmed by measurements. This result explains why the troublesome modes have high Qs because in the deformed cavity, the fields are shifted away from the HOM coupler where they can be damped as shown in Figure 1(b). The ideal and deformed cavities are shown in Figure 2. This analysis has shown that experimental diagnosis, advanced computing, and applied math worked together to solve a real-world problem as intended by the Scientific Discovery through Advanced Computing program. This example illustrates how powerful advanced simulation can be when applied to solving immediate problems.

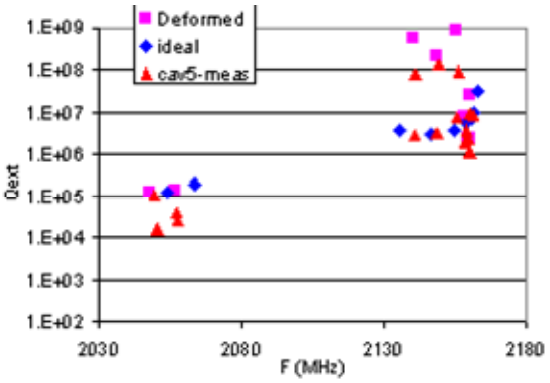

1 (a)

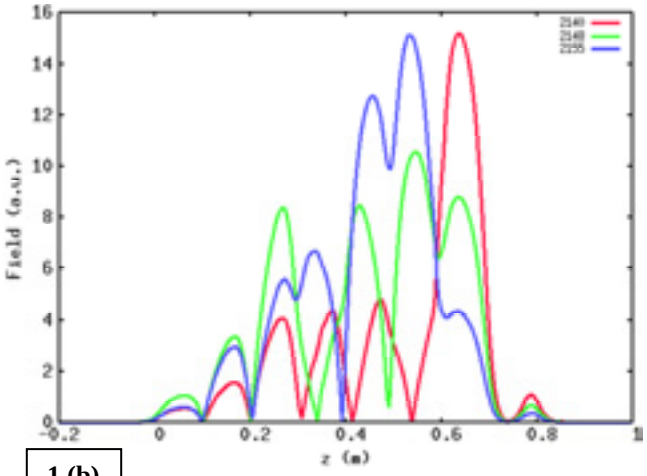

1 (b)

Figure 1 (a). Qext of calculated ideal and deformed cavities and measured values. **Figure 1(b)**. Field profile of high Q modes in deformed cavity. The damping HOM coupler is located at around z = -0.2 m, where the fields of these HOMs are very small. Images courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

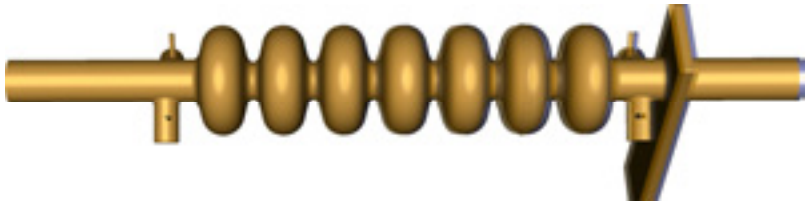

Figure 2. Ideal cavity in silver color and deformed cavity in gold. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

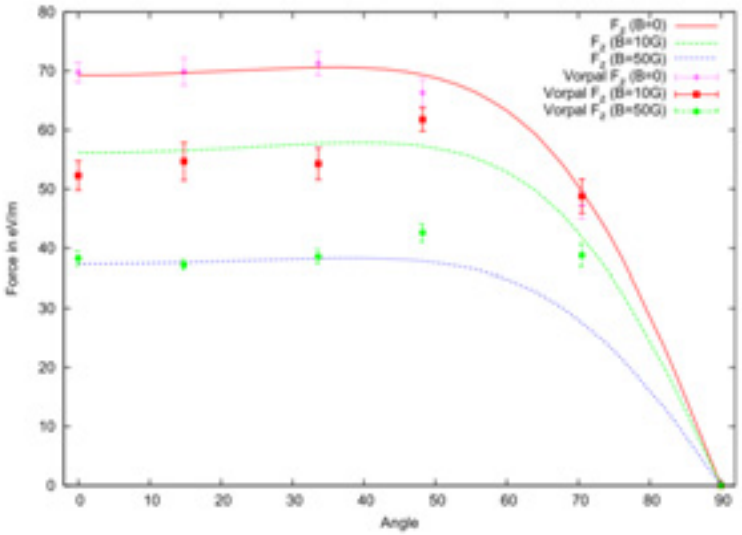

Figure 32. Parallel VORPAL simulations showing logarithmic reduction of the dynamic friction force in an undulator-based electron cooling system (Bell et al. 2008).

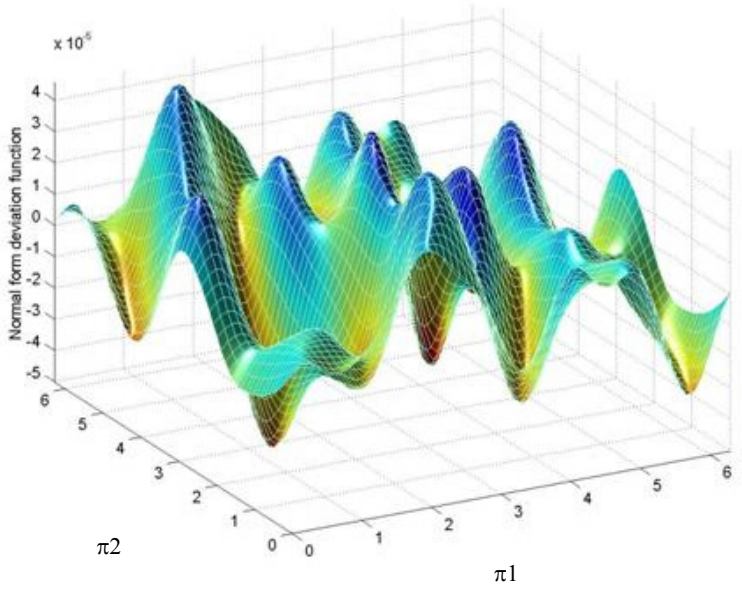

Figure 33. A magnified view of the normal form defect exhibiting a rich structure of local extrema at a minute scale in a two-dimensional subspace of the typically six-dimensional parameter space. Rigorous, tight bounding of the range of this function permits the generation of estimates regarding long-term stability of the dynamics. The problem is computationally challenging because of the small scale of the function; the large numbers of local extrema; and the need for tight, rigorous bounds. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Another application of such methods to a promising but complex synchrotron type is orbit stability studies for fixed-field alternating gradient accelerators (Prior 2007; Johnstone and Koscielniak 2008) (see Figure 34). The field shape is complicated, and as a result closed orbits and transfer maps around the orbits must be computed to a high order and in fine steps of up to 100 locations in the device.

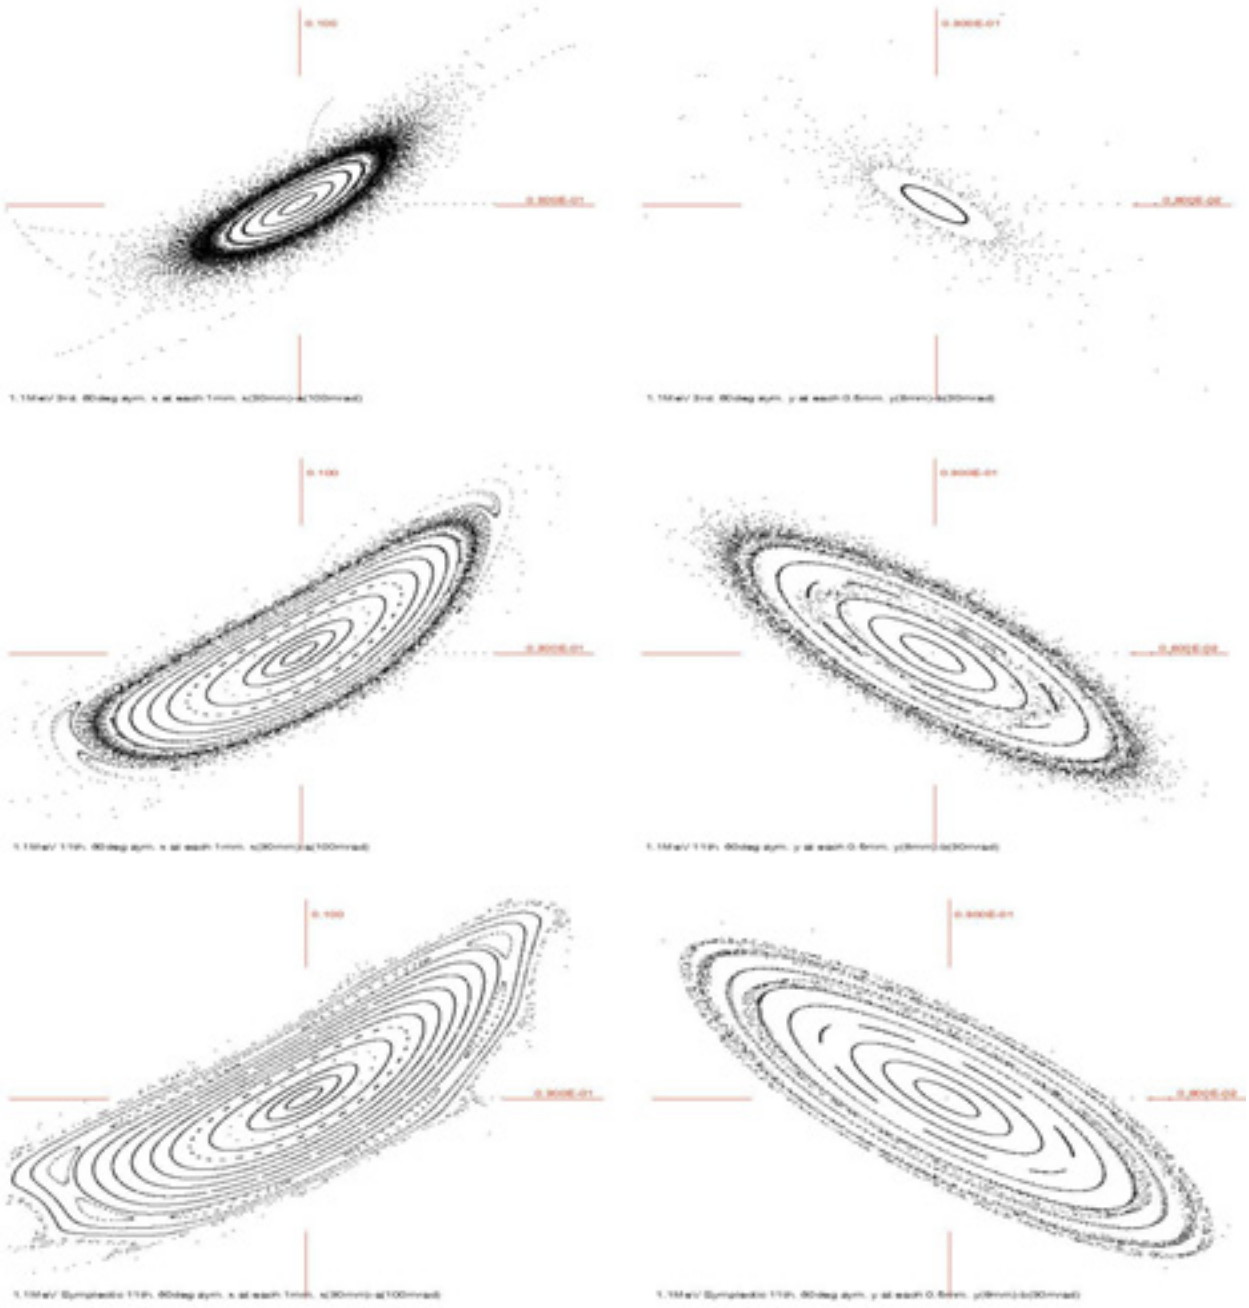

Figure 34. Long-term tracking of orbits in fixed-field alternating gradient type accelerator. Shown are horizontal phase space (left side) and vertical phase space (right side) modeled using transfer maps to third order (top), eleventh order (middle), and eleventh order with symplectification (bottom); even higher orders than these do not provide significant further changes. Due to the inherently nonlinear structure of the method, high orders are necessary for the simulation. Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

For each of these locations, it is common to perform a resonance analysis through normal form tools (Berz 1999) and orbit tracking of usually around 10^5 revolutions to assess stability. All these must be subjected to extensive design optimizations to arrive at viable machines. Currently, a sufficiently detailed field-design simulation and subsequent orbit analysis typically takes on the order of hundreds or

thousands of core-hours—and it is expected that an exhaustive automated search of parameter space may require on the order of 10^6 such iterations, leading to an overall cost in the range of hundreds of petaflop-years. This area is under active study and described further in the PRD section.

PRIORITY RESEARCH DIRECTIONS

Given the priorities of DOE's Nuclear Physics program and the opportunities afforded by the expected capabilities of extreme scale computing, the PRDs of nuclear physics accelerator science fall into four categories:

- research involving the proposed facility for FRIB
- research involving a proposed EIC
- research into optimizing complex electromagnetic structures for nuclear physics facilities
- research into advanced methods and applications of accelerator simulations.

Maximize Production Efficiency, Variety, and Purity of Rare Isotope Beams for Nuclear Physics Experiments

Basic Scientific and Computational Challenges

The design of FRIB poses several challenges. These include the development of new techniques for optimal design and tuning of isotope separators to select and to purify extremely rare isotopes, advances in electron cyclotron resonance (ECR) on source modeling, advances in radio-frequency quadrupole (RFQ) modeling, optimal design of low-beta radio frequency cavities, and advances in beam-dynamics modeling and optimization offline and in near real time. Figure 35 highlights several of the computational challenges associated with the design of FRIB, starting at the petascale and ranging to the extreme scale.

Optimal Design and Tuning of Isotope Separators

One of the most important aspects of experimenting with rare isotopes is related to the separation purity of the isotopes to be selected. If the reaction mechanism involves fragmentation and/or fission of heavy ion primary beams, then a multistage isotope separator is employed. The separator must be designed such that the primary beam rejection is perfect, and the several hundreds of other species that act as unwanted impurities for the experimental stations are minimized. The tiny production cross-sections of rare isotopes of interest to nuclear physics make this an especially difficult problem to model and simulate accurately and efficiently. The computational challenges are dependent on the rarity of the isotopes to be studied. These challenges include research currently possible on a typical single-processor personal computer (such as ^{14}Be obtained from fragmentation of oxygen requiring approximately10-gigaflop-days), cases that are manageable with petaflop machines (such as ^{132}Sn shown in Figure 35 requiring approximately 1 petaflop-day), and cases that require extreme scale computing resources (such as ^{100}Sn from fragmentation of xenon requiring approximately 100 petaflop-years).

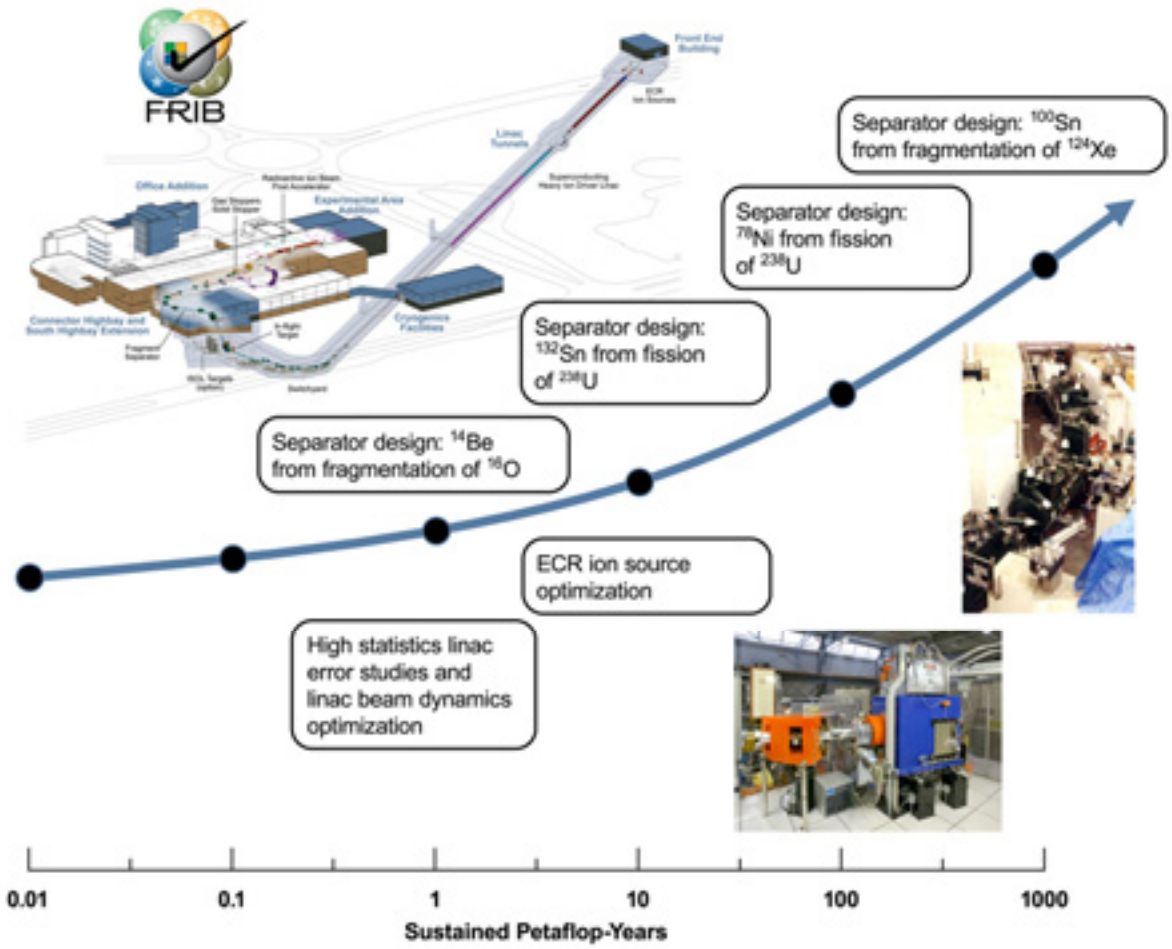

Figure 35. Anticipated highlights for priority research direction "Maximize Production Efficiency, Variety, and Purity of Rare Isotope Beams for Nuclear Physics Experiments." Upper-left image and right images courtesy of Michigan State University. Bottom-right image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

All cases noted above become extreme scale computing problems when design or experimental setup optimization is included. Global parallel-parameter optimization (magnet strengths, slit settings, target and wedge materials and thicknesses, and absorber shapes) is a challenge in its own, so the computational resource requirement is significantly increased with respect to that stated in the previous paragraph if these experimental parameters are added to the problem (factors of thousands more than above). These applications therefore need extreme scale computing resources. Some exceptionally difficult cases may require of the order of 100-exaflop-years.

The computational challenge may be even greater depending on the physics of the problem and the complexity of the parallel optimization model. Coupling to radiation transport codes, inclusion of space charge effects, and inclusion of full electromagnetic design in the optimization loop extend the problem to the extreme scale and beyond.

Advances in Electron Cyclotron Resonance Ion Source Modeling

For rare isotope facilities—such as the recently sited FRIB—ECR ion sources are required to produce a wide variety of highly ionized isotopes. These sources provide high currents of multiply charged ions for injection in the main driver accelerator. Large currents are needed for the physics program, and high charge states are needed to make most efficient use of available accelerating voltage.

Modeling ECR ion sources, including the ion beam extraction, is critical for better understanding and subsequent optimization of these sources. To date, reduced models have not agreed with experimental data. It is therefore necessary to use electrostatic and electromagnetic particle-in-cell (PIC) simulations that include multiple physical phenomena such as impact ionization, charge exchange, recombination, and particle-wall interactions. However, the extreme variation of spatial and temporal scales has limited such efforts to artificially small domains and short times (Mullowney et al. 2008). Emerging leadership-class computational facilities will enable full-scale, first-principle simulations.

The VENUS ECR source (Venus ECR) is ~1 m across, while the electron Debye length, λ_D, is ~10 microns, and the complicated magnetic fields (solenoid plus sextupole) are fully three-dimensional. The cell size can be chosen larger than λ_D if high-order particle shapes are used to prevent numerical heating. For $\Delta x \sim 50\lambda_D$, a mesh with $\sim 10^{10}$ cells is required, and a large fraction of the domain will have ~ 50 particles per cell (electrons plus multiple ion species), leading to $\sim 3 \times 10^{11}$ macroparticles. The timescale for ion dynamics (i.e., confinement, extraction, etc.) is several milliseconds, while the time step must be small enough to resolve the ionization/recombination timescale of 1 microsecond, which implies $\sim 10^5$ time steps.

For high-order particle shapes with multiple collisional processes, approximately 10 core-micro seconds are required for each particle step. Given a required number of 3×10^{16} particle steps in the calculation, and assuming efficient scaling up to approximately 1 million cores, this implies approximately 10^8 core-hour simulations, which translates into running at approximately 100 petaflop-hours. Given the need for a large number of such simulations (~1,000 per year) to design future and optimize existing ECR ion sources, the need for approximately 10 petaflop-years of computing resources can be predicted for each calendar year.

Advances in Electromagnetic Modeling and Optimization

Along with the above-mentioned advances in beam-dynamics modeling, advances in electromagnetic modeling and optimization that will be possible on extreme scale computers will have a major impact on the design of the FRIB. Examples include the design of ECR ion sources, of the FRIB RFQ, and of low-beta structures. These topics are described in more detail in the section titled, “Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities.”

Advances in Beam-Dynamics Modeling and Optimization

Currently, several highly parallelized beam-dynamics codes (e.g., TRACK and IMPACT) are available for the simulation of the FRIB accelerators. The majority of the accelerator design work can be performed using petascale computers, as demonstrated by teams from ANL, Michigan State University, and LBNL. The availability of extreme scale resources to support the commissioning and operation of the FRIB accelerators is essential to expedite the delivery, shorten the commissioning phase, and effectively operate the machine by applying a “model-driven accelerator” concept. Meanwhile, current parallel beam-dynamics codes need to be updated with large-scale optimization tools and must include

beam diagnostics and control systems into the optimization loop. These are described in the section titled, "Advanced Methods and Applications of Accelerator Simulation for Nuclear Physics Facilities."

Scientific Outcomes and Impacts

The availability of extreme scale computing resources, along with accelerator codes that make use of this computer power, will significantly impact the design and operation of FRIB, as well as other DOE SC accelerator facilities. In regard to FRIB, the availability of these resources will provide a means to design and operate isotope-separation systems with significantly increased accuracy and efficiency—and in some cases, will make seemingly impossible design simulations feasible. More generally, extreme scale computational resources will lead to optimized designs better able to meet facility requirements while reducing cost and risk. Beyond the design phase, extreme scale resources, applied in the model-driven accelerator paradigm, will permit faster commissioning, improved diagnosis of operational issues, and improved facility operations. The successful design, commissioning, and operation of the FRIB, facilitated by large-scale and extreme scale computing, will lead to numerous advances in nuclear science. These include providing a comprehensive description of nuclei, elucidation of the origin of the elements in the cosmos, providing an understanding of matter in the crust of a neutron star, and establishment of a scientific foundation for innovate applications of nuclear science (DOE 2007).

Develop Optimal Design for an Electron-Ion Collider

Basic Science Challenges and Computational Challenges

The design of a future EIC poses several challenges. To address these challenges, scientists focus on the development of advanced computational capabilities for multiphysics accelerator modeling, and the development of advanced electron-cooling systems and of energy-recovery linac technologies. Figure 36 highlights several of the computational challenges associated with the design of an EIC, starting at the terascale and ranging into the extreme scale.

Design of Electron Cooling Systems

Cooling of the hadron beam in a future EIC is critical to meet performance requirements needed for scientific discovery at such a facility. At present, electron cooling is the most promising concept. An electron cooling system has not yet been operated in the regime of a proposed EIC. Hence, accurate physical modeling is essential. Previous work based on molecular dynamics in a small, idealized domain needs to be extended to a molecular dynamics/PIC approach that covers the full extent of the overlapping electron and ion beams over many Debye lengths. This results in a multiscale simulation that requires extreme scale computational resources. For beam and system parameters relevant to the eRHIC concept, a single simulation using VORPAL operating with approximately 10% of peak performance requires approximately 10 teraflop-hours. To design an electron-cooling system, scientists must run approximately 10^4 seeds simultaneously to obtain enough dynamical friction force values to adequately characterize the performance for a single set of electron parameters. This implies a resource requirement of 100 petaflop-hours. These runs would need to be repeated for a variety of electron-beam parameters and also magnetic fields (solenoid and/or undulator) to optimize the design. For an optimal case, many instances would need to be run with different sets of magnetic field errors to evaluate the corresponding reduction in the cooling strength. Over the course of a year-long design effort, one could easily expect to complete approximately 1000 such cases, leading to resource usage of approximately 10 petaflop-years over the course of one calendar year.

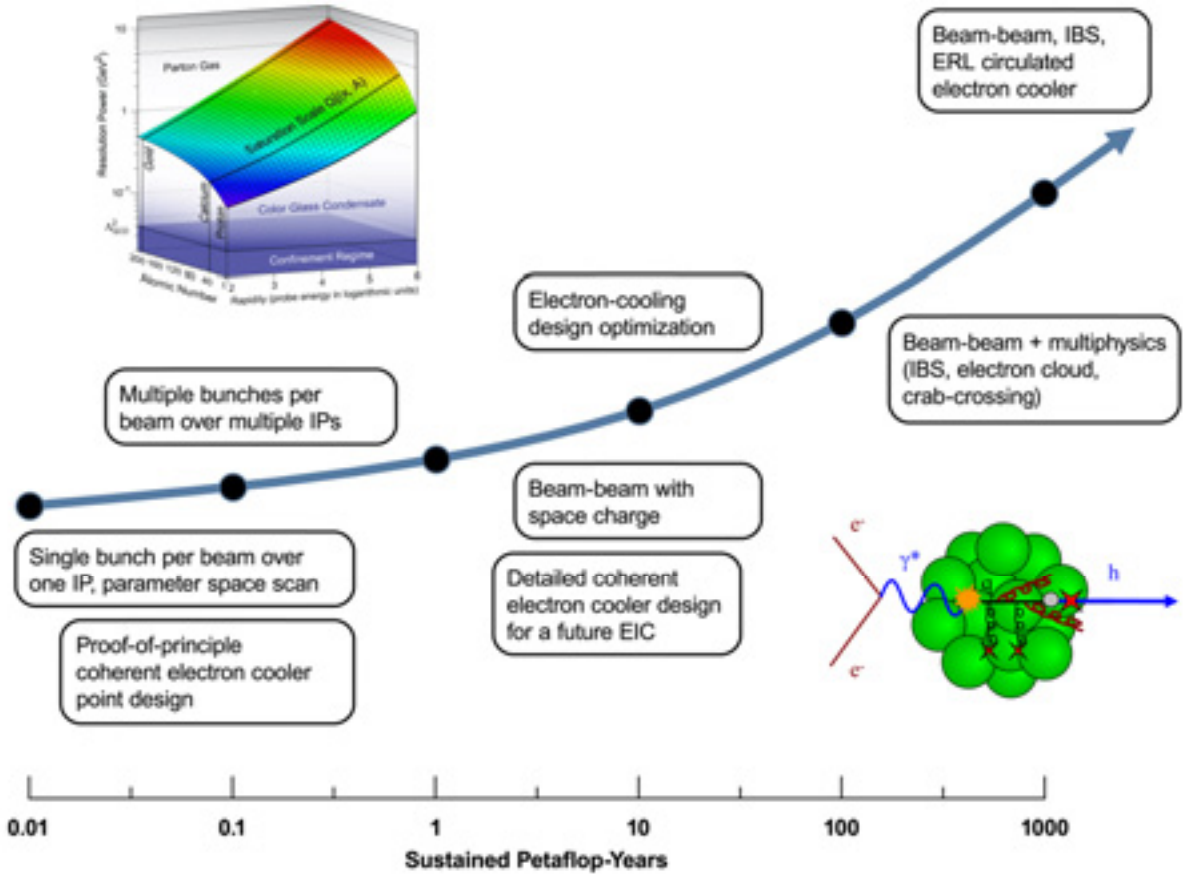

Figure 36. Anticipated highlights for priority research direction "Develop Optimal Design for an Electron-Ion Collider." Top-left image courtesy of Thomas Ullrich (Brookhaven National Laboratory). Bottom-right image courtesy of Alberto Accardi (Hampton University and Thomas Jefferson National Laboratory). Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Energy-Recovery Linac Technologies

The energy-recovery linac (ERL) is an important emerging concept for building energy efficient and therefore less expensive accelerators. Accelerator physics researchers are considering including ERLs in electron-cooling sections and EICs for nuclear physics projects. One important issue for understanding ERLs is beam loss due to beam halo. On the recovery pass, the beam phase space is significantly modified (e.g., due to energy extracted as part of the lasing or cooling processes); thus, nonlinear and nonideal behavior may result in increased beam loss. Modeling beams with sufficient computational particles to resolve beam halo and beam loss is a significant challenge that requires extreme scale computing.

A second important issue related to ERLs includes multipass beam breakup and coupling with higher-order modes. Again, because the beam on the recovery pass has a modified and nonideal phase space, the beam may couple more strongly with higher-order modes—especially at the higher currents considered for future machines. These modes may feed back and produce beam breakup. Modeling

beam breakup and coupling with higher-order modes will involve grid resolutions and simulation times that require extreme scale computing resources.

Finally, another important issue in ERL development is electron sources, where higher currents will produce issues of dark current, field emission, and cathode degradation due to ion bombardment. These higher currents will also produce subsequent heating issues associated with these effects. Modeling these effects in electron sources will require multiphysics code coupling and grid resolution that require extreme scale computing resources.

Scientific Outcomes and Impacts

The EIC will explore a new cold quantum chromodynamics frontier of strong color fields in nuclei, precisely image the gluons in the proton, and look for gluon saturation in heavier nuclei (DOE 2007). Large-scale, parallel computing is essential to design an EIC that maximizes luminosity and the potential for physics discovery. Beyond the development of an optimized EIC design, the design effort itself will lead to important new concepts and technologies for particle accelerators including new beam-cooling techniques, ERL technologies, advanced high-brightness sources, and new beam-beam compensation schemes.

Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities

Basic Science Challenges and Computational Challenges

Large-scale, high-accuracy electromagnetic design and optimization is essential to all major DOE SC accelerator projects. Several of the challenges and research directions have already been described in the "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale" workshop.[1] Examples include the development of novel, scalable eigen- and linear solvers for eigensystems at the extreme scale, as well as the implementation of efficient parallel algorithms for extreme scale modeling in a multiphysics environment. While many of the challenges and research directions are common to the various program offices of the DOE SC, each program office also has some unique needs and priorities. Evaluation of high-order mode (HOM) heating in the cryomodule of an ERL; design of the FRIB RFQ; and the multiphysics design and optimization of low beta radiofrequency accelerating cavities for FRIB—three priorities of the Nuclear Physics program—are described below. Figure 37 highlights the computational challenges associated with these areas, ranging from the terascale to the extreme scale.

High-Order Mode Heating in Cryomodule of Energy-Recovery Linac

One potential accelerator issue with ERL is HOM heating in the superconducting radio frequency cavities, especially with the combination of high current and short bunch length. It is important to estimate the broadband HOM power generated by the transit of the bunch to adequately dampen HOM without affecting cryogenic efficiency. To a great extent, the computational requirements are determined by the bunch length, whose frequency spectrum must be resolved by a fine mesh.

[1] "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale," December 9-11, 2008, Menlo Park, California. Additional information available at http://extremecomputing.labworks.org/highenergyphysics/index.stm.

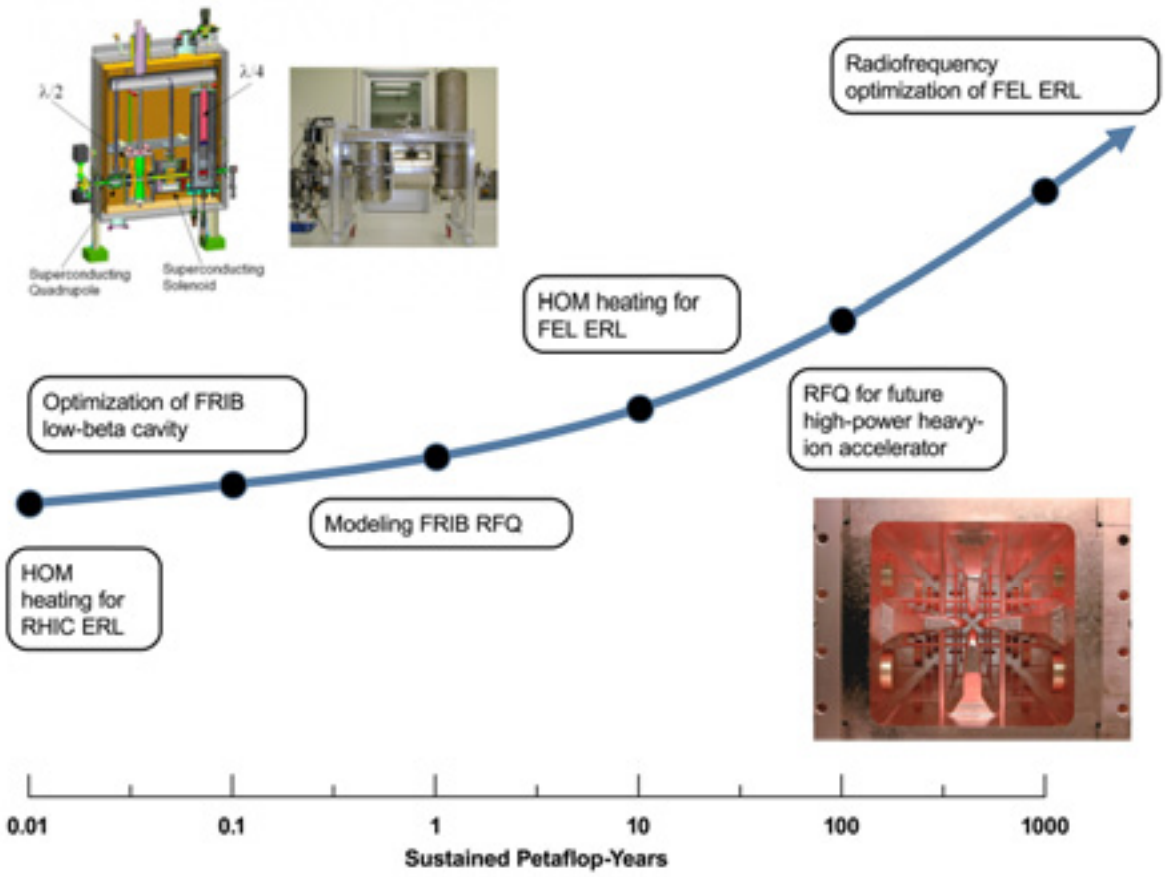

Figure 37. Anticipated highlights for priority research direction "Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities." Upper-left images courtesy of Michigan State University. Lower-right image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Thus, bunch length and mesh density are inversely related. In typical nuclear physics applications, such as the ERL used in the proposed luminosity upgrade of RHIC known as RHIC-II, the bunch length is relatively long (roughly 1 cm). The electromagnetic simulation of broadband HOM in a cryomodule involves solving a linear system with 150 million degrees of freedom for 100,000 time steps. The estimated computational requirement is approximately 1 petaflop-hour, which can be handled easily by a petascale computer.

In typical FEL applications, such as the ERL in the Jefferson Laboratory infrared FEL upgrade, the bunch length is much shorter (roughly 1 mm or less) than at RHIC-II. The electromagnetic simulation of broadband HOM in a cryomodule involves solving a linear system with 150 billion degrees of freedom for 10^6 time steps. The estimated computational requirement is approximately 1 petaflop-year.

For a cavity system on the order of 10 cm by 100 cm, a simulation with a resolution on the order of microns would require 10^{14} cells. Simulating many tens of oscillations will require 10^7 time steps, implying that 10^{21} cell steps are required. For electromagnetic PICs with a finite difference time domain

algorithm on current standard architecture, such as the Cray XT4[TM1], this implies 10^{14} - 10^{15} core-seconds. Therefore, to complete a simulation in a single day would require 10^{9} - 10^{10} cores, or a computing resource of approximately 3 petaflop-years.

Advances in Radio-Frequency Quadruple Modeling

The initial acceleration in FRIB is provided by an RFQ capable of accelerating any ion from hydrogen to uranium. The transport, bunching and focusing of charged particles in the RFQ resonator, is provided by appropriate design geometry of four modulated vanes. Parallel electromagnetic design tools with a fine-grained mesh must be used to develop a detailed resonator design for the RFQ. Accurate field maps of accelerating modes of cavities in the FRIB linear accelerator are required to reliably track particles in beam-dynamics studies. The long RFQ used in the low-energy linac could have disparate spatial scales ranging from fine transverse variations of the order of millimeters to longitudinal modulations of the order of a meter. While the frequency of the accelerator mode converges quickly with the reduction in finite element size, the field accuracy must be obtained with a much denser mesh or with the use of higher-order finite elements. The electromagnetic eigenmode simulation of the accelerating mode in the RFQ requires the solution of a linear system with 100 billion degrees of freedom. The estimated computational requirement is approximately 1 petaflop-day. This work is needed for future accelerators and accelerator upgrades, which will place greater demands on an injector than does FRIB, which will use the current state-of-the-art injectors.

Advances in Low-Beta Radio Frequency Cavity Design

The accelerator systems for the proposed FRIB will be based on several types of normal-conducting and superconducting transverse electric and magnetic mode (TEM-class) radio frequency cavities. The TEM-class superconducting cavities are much more efficient in the low-velocity (beta < 0.7) region compared to the elliptical superconducting cavities widely used for relativistic particles. Therefore, four to five types of TEM-class cavities are required in the FRIB driver linac. The integrated simulation of the resonators electromagnetic, thermal, and mechanical properties is necessary due to the complexity of the cavity's geometry. The optimal electrodynamic, thermal, and mechanical design of the fully dressed superconducting cavity together with fast tuner, slow tuner, radio-frequency coupler, and helium vessel could save millions of dollars. Obviously, the optimal design of the cavity should be free from multipacting. Currently, multiphysics design of a cavity can be performed using separate commercial software and does not allow researchers to provide full design optimization. Particularly, no code is available to study microphonics, which is due to coupled electro-mechanical oscillations induced by mechanical background noise. These simulations require a very large number of mesh points calling for extreme scale computing and a highly scalable multiphysics code.

Scientific Outcomes and Impacts

The development of an extreme scale, multiphysics cavity design and optimization tool will benefit not only the Nuclear Physics program but also the entire DOE accelerator complex. The process of developing these tools will also lead to advances in computational-enabling technologies such as linear solvers, eigensolvers, meshing, adaptive refinement, data analysis, and visualization.

[1] Cray XT4 is a trademark of Cray Inc.

Extreme scale electromagnetic modeling will shorten the design and build cycle for accelerator structures and components. It will also lead to cost savings by allowing researchers to explore and optimize designs that reduce cost while satisfying beam quality and machine operational reliability requirements. Lastly, the new modeling tools will provide accurate field maps of accelerating cavity modes for beam-dynamics simulation in overall machine studies.

Advanced Methods and Applications of Accelerator Simulations for Nuclear Physics Facilities

Basic Science Challenges and Computational Challenges

Extreme scale computing will dramatically change the way researchers design accelerators. It will enable the exploration of new accelerator concepts and the study of important phenomena, for which modeling was previously thought to be too computationally challenging. Extreme scale computing will also dramatically affect how scientists optimize accelerator designs. Lastly, extreme scale computing will change the way scientists' commission and operate accelerators.

Exploration of Advanced Concepts

Extreme scale computing will significantly impact the ability to explore innovative accelerators. The Accelerator Physics working group in the "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale" workshop identified the design of an ultracompact plasma-based collider as one of its key research directions.[1] The exploration and optimization of laser- and plasma-based concepts using extreme scale resources will have applications across the DOE SC. The design of fixed-field alternating gradient accelerators and the design of coherent electron cooling systems is discussed in the following sections.

The broad class of fixed-field alternating gradient accelerators is experiencing an international revival in the quest for high-beam power, duty cycle, reliability, and, in the case of the spiral-sector fixed-field alternating gradient, the potential for compactness at reasonable cost (Prior 2007; Johnstone and Koscielniak 2008, and references therein). The proposed fixed-field alternating gradients have the high average current and duty cycle characteristic of the cyclotron combined with the smaller aperture, beam losses, and energy variability of the synchrotron.

Because of frequently challenging field models, the computation of guiding, focusing, and accelerating fields and the assessment of stability of orbits usually must happen in an integrated mode of computation. Specifically, closed orbits and transfer maps around the orbits must be computed to a high order and in fine steps of up to 100 locations in the device. For each of these locations, it is common to perform a resonance analysis through normal form tools (Berz 1999) and orbit tracking of usually around 10^5 revolutions to assess stability. All these must be subjected to extensive design optimizations to arrive at viable machines. An example of one such simulation that illustrates the need for high-order models is shown in Figure 38. Currently, a sufficiently detailed field-design simulation and subsequent orbit analysis typically takes on the order of hundreds or thousands of core-hours—and it is expected that a

[1] "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale," December 9-11, 2008, Menlo Park, California. Additional information available at http://extremecomputing.labworks.org/highenergyphysics/index.stm.

future exhaustive automated search of parameter space may require on the order of 10^6 such iterations, leading to an overall cost in the range of exaflop-months.

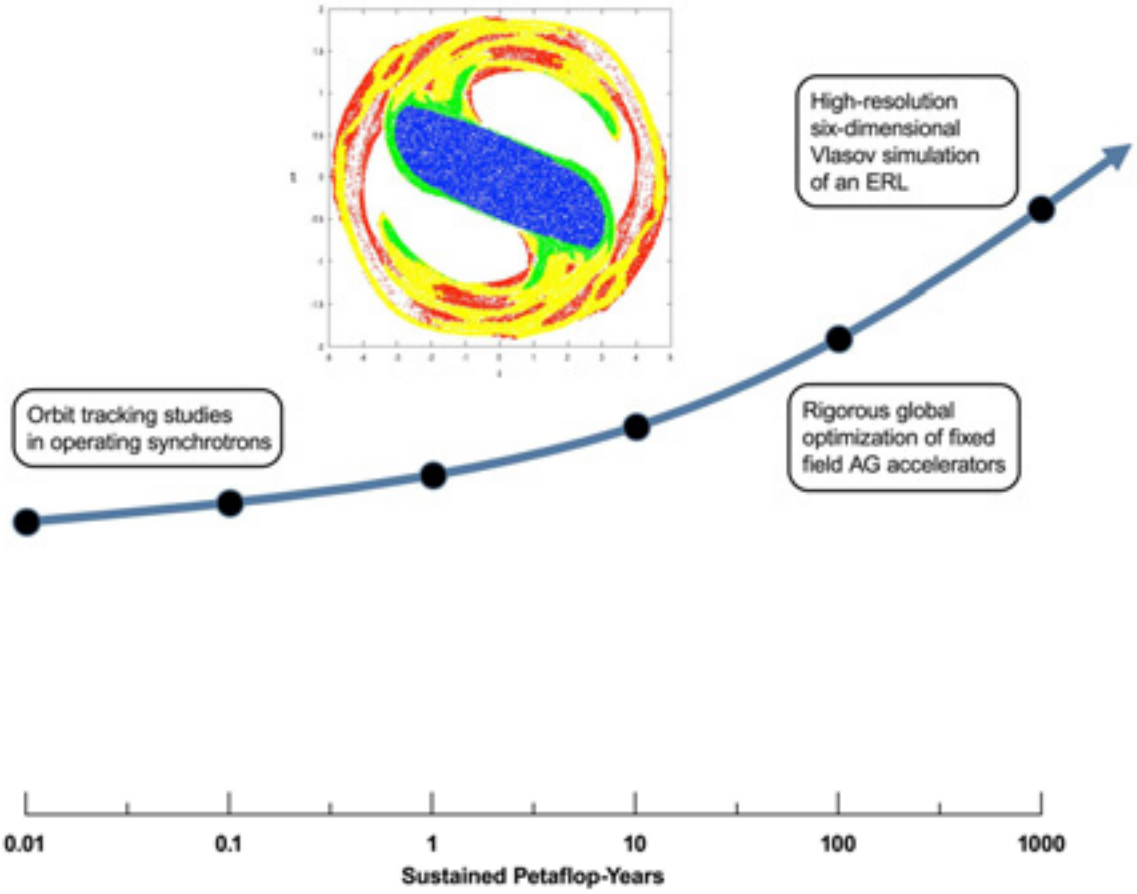

Figure 38. Anticipated highlights for priority research direction "Advanced Methods and Applications of Accelerator Simulations for Nuclear Physics Facilities." Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

The coherent electron-cooling concept proposes to combine the best features of electron cooling and stochastic cooling via free-electron laser technology. These features are combined to cool high-energy hadron beams on orders-of-magnitude shorter timescales (Litvinenko et al. 2008) than now possible. In a standard electron cooler, the key physical process is dynamic friction on the ions. The modulator section of a coherent cooler would be very similar to a standard cooler—but in this, case dynamical friction becomes irrelevant and the key physics is the shape of the density wake imprinted on the electron distribution by each ion. This implies use of a high-resolution PIC approach instead of the hybrid molecular dynamics/PIC approach used for conventional electron-cooling simulation. Though the high-resolution case involves approximately 1 billion cells, the time per step is comparable to the conventional case because resolutions of binary collisions are not required. As a result, the resource requirement is approximately the same as the conventional cooling scenario, namely approximately 10-petaflop-years over the course of 1 calendar year. The Poisson solver is a potential bottleneck, and it will be important that the solver for a single-point simulation scales to at least 10,000 computing cores.

Advanced Algorithms for Simulation and Optimization

The availability of extreme scale computing resources with massive memories will provide an opportunity to employ numerical algorithms that cannot realistically be implemented on current hardware. An example is the use of a direct Vlasov approach for the high-resolution prediction of ultra-low beam loss. Direct numerical simulation of the Vlasov equation has several advantages over the standard PIC approach. It allows for a direct and accurate description of particle distribution in the phase space; avoids random fluctuations and numerical noise caused by the finite number of macroparticles in a conventional PIC code; and enables scientists to accurately model low-density regions of the phase space such as the beam halo. Pure instability modes, which are difficult to study, are also conveniently simulated using PIC codes because the particle noise necessarily excites a spectrum of modes. Alternatively, the Vlasov approach is challenging computationally because calculations are performed in a (2n)-dimension phase space, where n = 1, 2, and 3; thus, the computational grid requires extreme computing resources. Current efforts have focused mainly on the two-dimensional (four-dimensional phase space) problem. To simulate multicomponent beams and plasma in a six-dimensional phase space, extreme scale computing will be necessary. For example, using 100 grid points in each direction requires 10^{12} grid points. The application of the Vlasov solver is an elegant approach to simulate three-dimensional problems in ion sources, be they electron cyclotron resonance or electron beam ion sources. These sources are inherent parts of the proposed FRIB.

Extreme scale computing resources will also provide the opportunity for employing new optimization algorithms. As an example, consider rigorous global optimization. Global parameter optimization will play a major role in the design of future nuclear physics particle accelerators. Recent advances in this rapidly progressing field allow the determination of rigorous solutions for global optimization problems with constraints. Based on dynamic domain decomposition and divide-and-conquer approaches, the methods are based on iteratively splitting the search space. On each of the currently active regions, a local or semilocal search is performed based on a variety of techniques including local descent-based method, linear or quadratic bounding, or genetic algorithms. More importantly, the currently active region is evaluated to determine whether, due to the behavior of the objective functions or the violation of constraints, it can be safely concluded not to contain the optimum, in which case it is discarded. While global parameter optimizations are very significant for the development of future particle accelerators, their computational requirements are immense. The underlying divide-and-conquer methods are currently performed on the largest available clusters and lend themselves to massive parallelization on 10^6 or more computing cores. One robust paradigm for parallelization is communication and load balancing on the "regular meeting" concept, at which all processors or suitable subgroups of processors share updates of the bounds of the objective function and redistribute large, unprocessed regions to achieve load balancing. Considering that a single evaluation of the objective function comprises a full simulation of the device in question—often comprising hundreds or thousands of core-hours—and that an exhaustive search can involve thousands to, in extreme cases, millions of evaluations of the objective function, tasks may result that require exaflop-months.

Advanced Tools for Accelerator Commissioning, Operation, and Control

Today, no accelerator facility in the world can fully rely on a computer model for its operation. There are intensive efforts to construct the needed modeling environment using experience gained at current state-of-the-art accelerators such as the Spallation Neutron Source, RHIC, and Jefferson Laboratory. A promising approach is a configurable framework, independent of any specific accelerator that could allow

insertion of computing modules for different scale simulation, optimization, and control algorithms. Careful development of portable and mutable software types, or classes, to describe the accelerator elements is part of this and is another area for fruitful collaboration with applied mathematicians and computer scientists (Malitsky et al. 2009).[1] As more intense and complex machines are built, the challenge of achieving their design goals and operating them effectively grows larger. Using traditional operating and beam-tuning methods (mostly manual) will certainly result in commissioning delays and reduced machine availability for science. A complex accelerator system, as is being proposed for FRIB, uses primary beams ranging from proton beams to uranium beams, at different energies to produce secondary beams of rare isotopes spanning the whole chart of nuclides. Such an accelerator will require currently unavailable advanced tools to support operations. Using a parallel, realistic three-dimensional beam dynamics model to support commissioning and real-time machine operation will expedite the delivery of the machine, greatly enhance its availability, and significantly reduce its operating budget. Supporting online machine operations involves large-scale optimization problems; fast interfaces between the computer model, the beam diagnostic devices, and the beam line elements; and advanced data-analysis and -visualization tools. Combined with the requirement of fast turn-around calculations (seconds to minutes) to support fast decision making, the task can easily reach the extreme scale. This concept of a model-driven accelerator will benefit existing—and more importantly future—nuclear physics facilities.

Scientific Outcomes and Impacts

The development of advanced methods for extreme scale modeling of nuclear physics accelerators, and the application of these methods to novel types of accelerators, will have a major impact on future nuclear physics accelerator facilities. It will enable the exploration and possible usage of novel types of accelerators and accelerator systems (such as fixed-field alternating gradients, new methods of beam cooling, etc.). These methods will also lead to advanced design optimization techniques that will also benefit other types of DOE SC projects and lead to new tools for real-time control of large scientific experiments and facilities.

APPLIED MATHEMATICS AND COMPUTATIONAL ISSUES

Particle accelerator science and technology spans most of the offices within DOE SC. The Nuclear Physics, High Energy Physics, and Basic Energy Sciences programs all develop and operate major accelerator facilities. In addition, the Fusion Energy Sciences program has accelerator-based programs related to inertial fusion and high-energy density physics.

With regard to the Nuclear Physics, High Energy Physics, and Basic Energy Sciences programs, the primary accelerator technology is radio frequency technology, and the modeling issues strongly overlap. The “Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale” workshop identified a number of applied mathematics and computer science research directions that will be relevant to future high-energy physics accelerator-modeling activities at an extreme

[1] Wang N. “Next Generation Communication Middleware for High Level Accelerator Control Systems.” Paper presented at the *2009 International Computational Accelerator Physics Conference* in San Francisco, California. Not publicly available.

scale.[1] Those research directions are essentially the same for the Nuclear Physics program and the accelerator modeling managed by Basic Energy Sciences. Examples include sparse matrix algorithms, meshing, large-scale partial-differential-equation-constrained optimization, multiobjective optimization, load balancing, massive data analysis, framework development, input/output optimization, single-node performance, and fault tolerance.

CONCLUSIONS

High-performance computing and modeling is essential to the success of future nuclear physics accelerator facilities such as FRIB and EIC; in fact, it is essential to the design of all major DOE accelerator facilities. Advanced simulation allows exploration of designs and concepts in a virtual environment that would otherwise be prohibitively expensive or otherwise impossible to explore in the real world. Through computer-aided parallel optimization, it is possible to perform realistic simulation and design optimization to achieve accelerator design objectives for accelerator components, systems, and facilities as well as to minimize cost and risk. Development of a new generation of computational resources that significantly exceed present-day petascale capabilities—as well as new computational capabilities for extreme scale modeling of accelerator science and technology—will allow this improved design of new accelerators to take place.

Extreme scale computing resources may revolutionize accelerator design, commissioning, and operation. However, using these resources effectively will require advances in computer hardware, as well as in a number of supporting areas. A strong program of research and development in supporting technologies including applied mathematics; computer science; parallel algorithms, optimization, and numerical libraries; supporting software; data storage and analysis; visualization; and visual analytics will allow the accelerator work to take advantage of any advanced computing hardware deployed.

The accelerator activities of the various DOE SC program offices have much in common. For example, both the Nuclear Physics and High Energy Physics programs require the ability to accurately design and optimize colliders and beam cooling systems. The Nuclear Physics, High Energy Physics, and Basic Energy Sciences programs require the ability to model high-brightness electron sources and ERLs. Projects sponsored by DOE SC program offices are required to have the ability to accurately model space-charge effects and geometrically complex three-dimensional electromagnetic structures.

Accelerator activities sponsored by the DOE SC program offices also have differences. For example, multicharge state transport is mainly important to the Nuclear Physics program. Accurate modeling of coherent synchrotron radiation is mainly important to the Basic Energy Sciences. Modeling of neutralized transport is mainly important to Fusion Energy Sciences. Nonetheless, it is clear the commonality far outweighs the differences. Broad efforts in accelerator modeling can simultaneously address both the issues common to all these areas, as well as the unique needs of specific design and science program areas.

[1] "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale," December 9-11, 2008, Menlo Park, California. Additional information available at http://extremecomputing.labworks.org/highenergyphysics/index.stm.

Table 5 provides an outline of the milestones for the work described in this section. Provided that the computational resources are made available for research in accelerator physics at the anticipated scales, the forefront research that will be conducted are provided as milestones.

Table 5. Milestones for Accelerators Physics

Scale	Milestone
>1 Petaflop-year	• Separator design for ^{14}Be from fragmentation of ^{16}O • Linac error studies and linac beam dynamics optimization • Proof-of-principle coherent electron cooling channel design • Parameter-space scans for multiple beam-bunches for multiple collider interaction points • Modeling of RFQ for FRIB
> 20 Petaflop-years	• Separator design for ^{132}Sn from fission of ^{238}U • ECR ion source optimization • Electron-cooling design optimization • Beam-beam space charge • Detailed coherent electron cooling channel design for an EIC • Higher-order-mode heating for an ERL for a free-electron laser
> 100 Petaflop-years	• Separator design for ^{78}Ni from fission of ^{238}U • Optimization of an advanced RFQ for future high-power heavy-ion accelerators
>1 Exaflop-year	• Separator design for ^{100}Sn from fragmentation of ^{124}Xe • Beam-beam, intrabeam scattering for an ERL circulated electron cooler • Beam-beam and multiprocess physics together, including intrabeam scattering, electron-cloud feedback, and crab-type crossings • RF optimization for an FEL-based ERL • Fixed-field alternating gradient accelerator optimization • Six-dimensional Vlasov for an ERL

CROSSCUTTING CHALLENGES

CROSSCUTTING CHALLENGES

From a general scientific perspective, many overlaps exist between nuclear physics and other fields. For example, several common problems are addressed by high-energy and nuclear physicists, who often employ similar theoretical and experimental techniques to understand physical processes of interest. These problems include exploring the nature of the electroweak interaction and physics beyond the standard model, studying how the strong interaction arises from quantum chromodynamics (QCD), and especially investigating how quarks are confined into observable particles. Similarly, the connections between the nuclear physics and astrophysics communities are multifaceted and deep. Those connections have developed into the common field of nuclear astrophysics, which investigates many phenomena such as the following:

- stellar evolution
- origin of the elements
- Big Bang cosmology
- solar physics
- neutron star and pulsar physics
- explosive phenomena such as supernovae, gamma-ray bursts, novae, X-ray bursts, and cosmic-ray physics.

To address these issues, nuclear physicists, astrophysicists, and astronomers must work together to observe, describe, and calculate astrophysical phenomena. A similar coherence exists in quantum many-body problems in chemistry, materials sciences, biology, and nuclear structure—where theorists solve the same Schrodinger equation with similar techniques with the only difference being the interaction between particles. Basic research into the physics of the nucleus, coupled with research into the chemical properties of radioactive elements, has established a solid foundation for practical technologies such as nuclear energy, nuclear medicine, medical imaging, particle accelerators, and particle and radiation detectors. Moreover, in the course of this basic research, nuclear scientists have created a host of tools and instruments that have proved valuable in the marketplace such as radioisotope generators used in nuclear medicine and satellite power sources. They have compiled a wealth of essential data about nuclei and the many reactions in which nuclei engage. Furthermore, nuclear scientists have helped train, and continue to train, the highly specialized workforce needed to sustain and advance these applied nuclear technologies.

From a computational perspective, nuclear physics spans a breadth of efforts focused on determining the properties and dynamics of matter under extreme conditions, nuclei with extreme proton-neutron ratios, and explosive stellar dynamics. The computational nuclear physics field overlaps in certain key areas. For example, lattice quantum chromodynamics (LQCD) methods are used to investigate the nature of deconfinement in the quark-gluon plasma (hot QCD); they are also used to derive from first principles the confinement of quarks and gluons into protons, neutrons, and a broad range of hadronic matter (cold QCD). A growing research area within LQCD involves computing the interaction between protons and neutrons. This connection will lead to a description of nuclear forces developed directly from QCD, a long-sought goal. These interactions will feed directly into efforts to describe nuclei—from light to medium mass, to neutron rich, and on up to super heavy elements. Furthermore, research in fundamental symmetries will help quantify the neutron electric dipole moment arising from the potential charge-

conjugation and parity-violation within QCD, which has relevance to high-energy physics and extensions of the standard model to address—for instance, the observed matter-antimatter imbalance in the universe. Nuclei and their reactions are responsible for element production in the universe and for the evolution of and power generation in stars. Neutrinos that power the explosions of supernovae are generated through interactions with protons and nuclei in the center of the star, and also originate from the p-p reaction that is both the main power generation step in most of the life of the star and the first step in stellar nucleosynthesis. Further, nuclei are used in laboratories to discover fundamental neutrino properties.

This document notes several key developments that will enable nuclear physics and nuclear astrophysics to advance during the evolution to extreme scale computing.

- Computational algorithms and codes must exhibit the scaling necessary to run efficiently on platforms with millions of computing cores and on platforms that may exploit accelerators. Efficient use of memory, communications among nodes, and combinations of shared and distributed computing will be required to maximize the potential of the next generation computers. This will require exceptional work on the front end to develop scalable codes that can accomplish the science.
- Load balancing of the computation, particularly for simulations that deploy adaptive mesh refinement and quantum Monte Carlo algorithms, will remain important. Additionally, simulation fault tolerance as the number of computing cores used in a given calculation progresses into the millions, and beyond, will become even more important than today.
- Because the amount of memory per computing core will not increase—and it may decrease as scientists approach the extreme scale—development of new algorithms for efficient domain decomposition and load balancing will be required.
- New parallel input/output (I/O) algorithms that can handle files of many terabytes in size and beyond, along with mass storage that can accommodate exabytes of data, must be developed. For example, the turbulent nature of the deflagration phase of an exploding supernova demands high temporal resolution in the retained data sets. This leads to the production of remarkable data volumes (many petabytes and even exabytes). Similarly, QCD field configurations at decreasing lattice spacing and nuclear many-body wave functions both require extreme amounts of I/O to store and retrieve results that must be computed and stored for later use and analysis.
- The growing volume of data, with an increasingly ambitious physics program, requires sufficient computational resources for post-processing of the data. This will entail the provision of computer systems that are themselves large scale by current standards, and with an aggregate capacity of at least the scale of the extreme (capability) resources themselves. Thus, the enterprise of computing will require an “ecosystems” level approach to staging, executing, and post-processing data that come from extreme scale computations.
- New algorithms for scientific data analysis, including visualization, must be developed to handle petabytes and exabytes of data. Other algorithms to be developed include data-archiving techniques that can process exabytes of data and allow for comparative analyses to be performed between huge data sets.
- An efficient, collective parallel I/O from millions of computing cores must be established. Data management approaches for geographically distributed teams and data analysis algorithms must be developed for what will ultimately be petabytes of data per simulation delivered over the course of days to months. Discovery-enabling visualization of multivariate (scalar, vector, and tensor), multidimensional (as high as six-dimensional), petascale data must be developed.

- Because the final products of LQCD calculations are hadronic correlation functions, a reconsideration of the methods of their calculation from quark-propagators is required before moving to extreme scale computing. The need for extended precision arithmetic has recently been demonstrated across several problems in QCD.

While the above points are emphasized in this report, investigations of nuclei and nuclear astrophysics can benefit other fields and will be informed by advances in those fields. For example, the complex, multiphysics simulations that will be required to develop precision stellar and supernovae models will yield computational methodologies and tools that can benefit a number of other critical application domains including climate modeling, simulations of fusion energy devices, simulations aimed at the development of new and more-efficient combustion engines, and stockpile stewardship. In particular, computational astrophysics research will lead to effective computational approaches for simulating turbulent fluid flow, for radiation transport, and for radiation hydrodynamics. Radiation transport—particularly photons and neutrinos, which can couple a problem over long scales—emerged as a key crosscutting area.

In summary, computational techniques and needs complement the scientific areas that will be pursued with extreme scale computing. Examples include, but are not limited to, the following:

- improved linear algebra techniques for large matrices
- massive global nonlinear optimization with nonlinear constraints
- fault tolerance
- asynchronous I/O and load balancing across hundreds of thousands to millions of cores
- improved programming environments
- verification and validation issues for extreme scale computations.

Managing these efforts requires methodologies for meeting the challenges associated with scientific simulation workflows, including data management and analysis, visualization, workflow management, and automation. The scientific community views these as important problems to address as progress in computational hardware continues.

To conduct the broad range of computing at the extreme scale in nuclear physics and nuclear astrophysics, investigation and implementation is needed of a funding model that enables domain-specific scientists—i.e., computational nuclear theorists—applied mathematicians, and computer scientists to work together on specific problems. For the nuclear and astrophysics scientific communities, the U.S. Department of Energy's (DOE) Scientific Discovery through Advanced Computing model currently serves this need as DOE progresses toward the petascale. The scientific community supports an expanded version of Scientific Discovery through Advanced Computing as part of DOE's overall plan to move to extreme scale computing. Algorithmic developments for extreme scale machines will also entail having advanced knowledge of hardware changes associated with new machines. Early access to such information (through nondisclosure agreements) would be valuable in planning for and developing algorithms that can use the computational resources immediately upon deployment.

PRIORITY RESEARCH DIRECTIONS

COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES

NUCLEAR STRUCTURE AND NUCLEAR REACTIONS

NUCLEAR ASTROPHYSICS

HOT AND DENSE QUANTUM CHROMODYNAMICS

ACCELERATOR PHYSICS

COLD QUANTUM CHROMODYNAMICS AND NUCLEAR FORCES

SPECTRUM OF QUANTUM CHROMODYNAMICS

Computing the bound state spectrum of quantum chromodynamics (QCD) is vital if scientists are to claim a complete description of the strong interactions, and the confrontation of high-precision calculations of the spectrum with future experimental measurements is a vital test of the theoretical framework. Figure 39 illustrates significant milestones that will impact scientists' understanding of QCD as available computing resources evolve toward the extreme scale. In contrast to electromagnetism, the "field-lines" between a quark and antiquark in QCD do not diffuse over large distances, but rather are confined to compact "flux tubes" connecting the quark and antiquark. Baryons themselves are emblematic of QCD, with the three quarks carrying each of the three color charges of QCD. The outstanding arena that spectroscopy provides for exploring QCD is driving intense experimental studies of the spectrum, primarily excitations of the "glue," or gluonic degrees of freedom, with the GlueX experiment, a flagship component of the 12 GeV upgrade at the Thomas Jefferson National Accelerator Facility (Jefferson Laboratory). Extreme computing will provide the *ab initio* theoretical calculations required to capitalize on these experimental investments.

Basic Scientific and Computational Challenges

Unstable Resonances. Lattice quantum chromodynamics (LQCD) provides an *ab initio* method to compute the meson and baryon spectrum by exploiting the finite spatial extent of the gauge-field configurations that are used in such calculations (a "finite volume"). In particular, variation of the lattice spatial-volume provides a mechanism to compute the scattering phase shifts for the resonances and their decay modes. A challenge in the approach to the extreme computing era is the extension of currently established methods for investigating elastic processes to the treatment of inelastic decays, in which there are multiple final states.

Flavor-Singlet Contributions to the Spectrum. Calculations of the spectrum have largely been confined to systems that do not admit the annihilation of an initial-state quark with an initial-state antiquark. The inclusion of such terms will require the calculation of so-called "disconnected contributions" (or "disconnected diagrams") with sufficient precision that the energy spectrum can be resolved. This will necessitate the introduction of improved stochastic estimators or the development of alternative methods.

Improved Statistical Analysis. As the energy of a state is increased, and as the light-quark masses are decreased, the signal-to-noise ratios of the correlation functions associated with these states generally degrade severely. Current methods provide powerful tools for delineating the different states, but their effective use will require the development of improved statistical tools to fully exploit the investment in leading-edge computing.

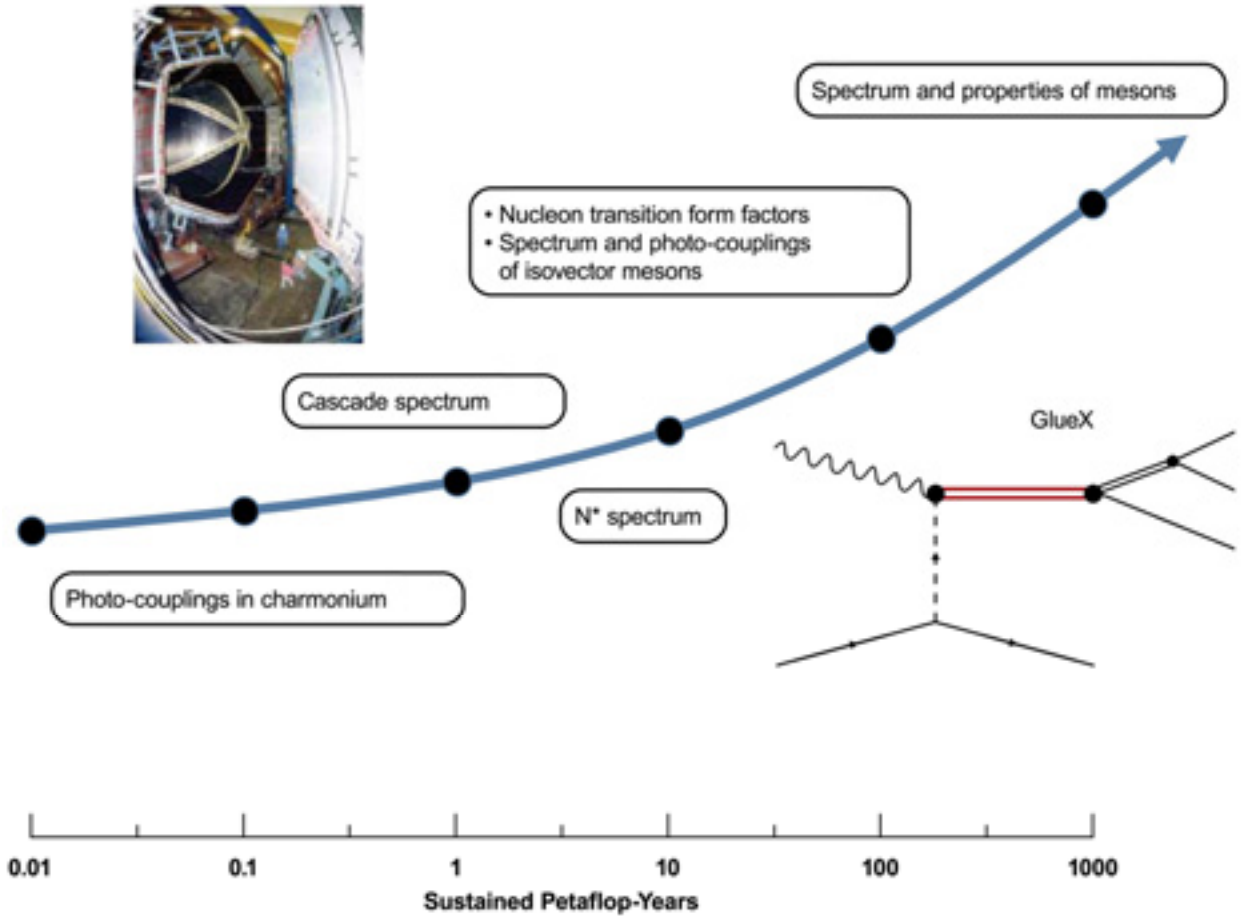

Figure 39. Anticipated highlights for priority research direction "Spectrum of Quantum Chromodynamics." Upper-left image courtesy of the Thomas Jefferson National Accelerator Facility. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Scientific Outcomes and Impacts

The Spectrum and Properties of Meson Resonances. The presently observed spectrum of QCD provides little direct evidence of the presence of gluons. However, QCD presents the possibility of exotic mesonic states of matter in which the gluonic degrees of freedom are explicitly exhibited, and the flux tubes excited. The search for such states will be the subject of intense experimental effort, notably the GlueX experiment at the 12 GeV upgrade at the Jefferson Laboratory (JLab@12GeV). The confrontation of the precise LQCD calculation of the spectrum afforded through extreme computing with the experimentally determined spectrum of meson resonances will provide the culmination of the quest to understand QCD as the theory of strong interactions. The calculation of the spectrum and properties of exotic resonances will reveal the nature of the gluonic degrees of freedom in the spectrum, and may help elucidate scientists' understanding of the origin of confinement.

The masses and widths of cascade resonances, analogues of the proton and neutron but with two of the *u* and *d* quarks replaced with the heavier, strange quarks, are poorly determined. Even the quantum numbers of many of these states are unknown. Their decay widths are expected to be small and their investigation in LQCD correspondingly less demanding. Computation of the cascade spectrum will require approximately one petaflop-year, and should provide clues as to the role of quark flavor and mass in the spectrum of QCD. Further, these computations are another opportunity for LQCD to provide predictions for future experimental searches.

The spectrum of N* resonances is the subject of intense experimental activity, with its importance encapsulated in the DOE 2009 milestone HP3 (DOE 2008):

> Complete the combined analysis of available data on single π, η, and K photo-production of nucleon resonances and incorporate the analysis of two-pion final states into the coupled-channel analysis of resonances.

The computational methods developed to determine the spectrum of cascades can be extended, but the greater range of decays makes this a more challenging computation, requiring tens of petaflop-years. The baryon spectrum is emblematic of the non-Abelian nature of QCD, and key questions being addressed include the following: what are the roles of the gluons, and more specifically, the role of gluon self-interactions in nucleons? More generally, what are the effective degrees of freedom describing the baryon spectrum?

The experimental measurement of the electromagnetic transitions between low-lying N* resonances are encapsulated in the DOE 2012 milestone HP7 (DOE 2008):

> Measure the electromagnetic excitations of low-lying baryon states (< 2 GeV) and their transition form factors over the range $Q^2 = 0.1 – 7$ GeV^2 and measure the electro- and photo-production of final states with one and two pseudoscalar mesons.

The LQCD calculation of these transitions will require approximately 100 petaflop-years, and provide further clues to the composition of the low-lying baryon spectrum. Furthermore, calculation of the electromagnetic properties with increasing Q^2 (square of the four-momentum transferred to the hadron) enables the perturbative QCD approach to a quark and gluon picture of hadrons to be investigated.

The future GlueX experiment at Jefferson Laboratory's 12 GeV upgrade aims to photo-produce so-called exotic mesons, with the first physics results expected in the middle of the next decade. LQCD has a vital role in both predicting some of the low-lying spectrum, notably for those states with isovector quantum numbers, but also in computing the photo-couplings between these and conventional mesons. These calculations will provide vital input for estimating production rates in the GlueX experiment, and highlights the role of LQCD in guiding experiments.

HOW QUANTUM CHROMODYNAMICS MAKES A PROTON

Protons and neutrons, collectively known as nucleons, are the basic building blocks from which all nuclei are constructed, but are themselves formed from the quarks and gluons of QCD. Determining how the quarks and gluons form protons, neutrons, and other hadrons is at the core of frontier nuclear physics experiments at the Brookhaven National Laboratory in New York, the Jefferson Laboratory in Virginia, and international laboratories. Extreme computing is required to perform *ab initio* LQCD calculations of

the fundamental properties of nucleons, and provide insight into their structure that is inaccessible to experiment. Together, forefront LQCD calculations and new experimental measurements—such as those exploring transversity and of generalized parton distributions—will enable scientists to build a three-dimensional picture of neutrons and protons in terms of the primordial quarks and gluons of QCD. Finally, scientists will discern how mass, spin, charge and currents are distributed within a nucleon.

Basic Scientific and Computational Challenges

Calculation of Gluon Contributions to Hadron Structure. Although approximately half the momentum and spin of the nucleon comes from "glue," or gluonic degrees of freedom, calculations of the gluonic contributions within hadrons are far more difficult than those of the corresponding quark contributions. Improved gluonic operators must be developed and the computational infrastructure for much higher statistics calculations will be needed.

Calculation of Flavor-Singlet Gluon and Sea-Quark Contributions. Precision calculations of hadron structure have been largely restricted to *isovector* quantities, such as the difference between the proton and the neutron matrix elements, in which the so-called disconnected contributions cancel and gluons do not contribute. Calculations of proton and neutron properties separately, and more generally the flavor-separated contributions of quarks and gluons to hadron structure, require calculation of disconnected diagrams and their mixing with gluons. Practical calculation of these notoriously difficult quantities will require the development of improved estimators and stochastic noise techniques.

Higher Moments of Structure Functions. Because the ultimate goal is to calculate structure functions and LQCD calculations can only produce moments of these functions, it is desirable to calculate as many moments as possible to optimally reconstruct the relevant physics. Because the lattice has hypercubic symmetry, and not the Lorentz symmetry of the space-time, present techniques only permit calculation of the three lowest moments of the structure functions. Thus, it is necessary to develop new techniques to enable calculation of higher moments.

Form Factors at High Q^2. The ability to determine hadron structure at very short distances, or alternatively at high-momentum transfers, is limited by systematic uncertainties associated with the finite lattice spacing in LQCD calculations, by the degrading signal-to-noise ratios at increasing hadronic energies, and by the decreasing size of the form factors at high momentum.

Figure 40 illustrates the progression toward extreme computing for hadron structure.

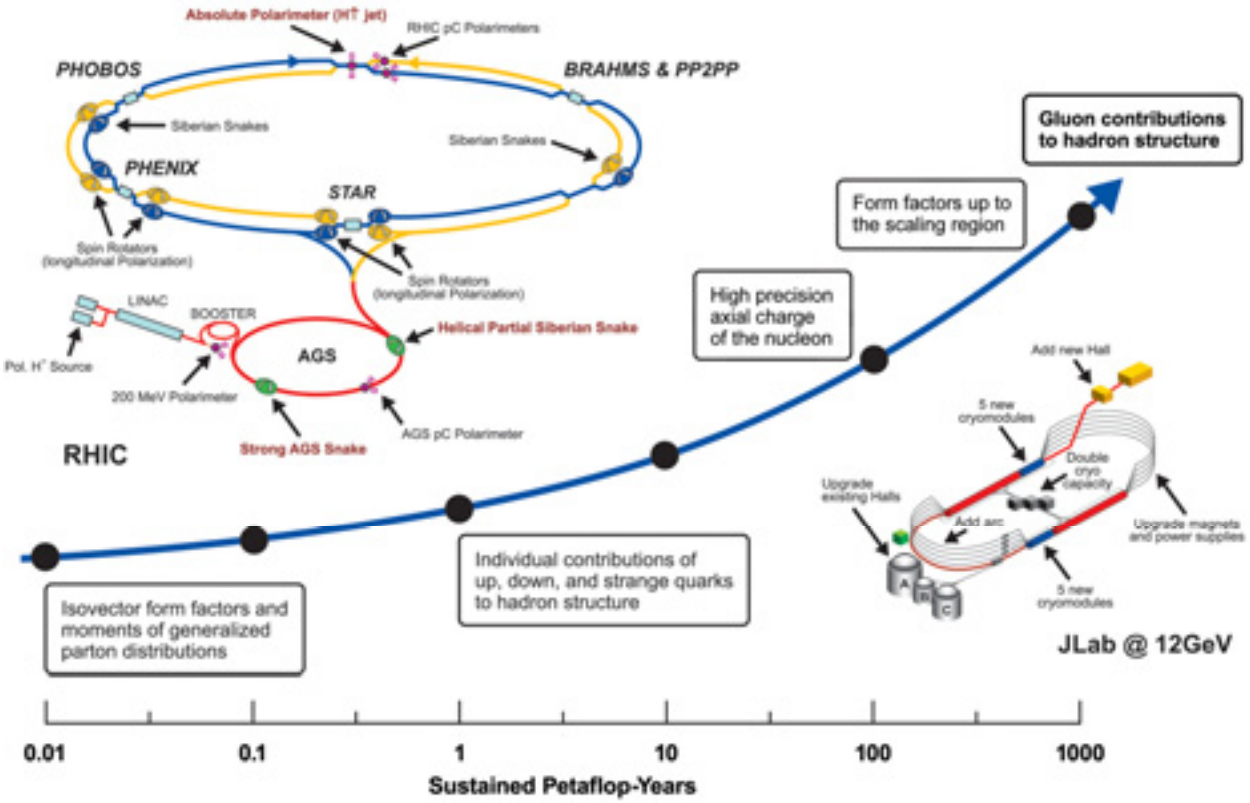

Figure 40. Anticipated highlights for priority research direction "How Quantum Chromodynamics Makes a Proton." Upper-left image courtesy of Brookhaven National Laboratory. Lower-right image courtesy of Thomas Jefferson National Accelerator Facility. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

Scientific Outcomes and Impacts

Gluon Contributions to Nucleon Structure. The contribution of gluons to the nucleon mass, and the calculation of the low moments of the spin-averaged and spin-dependent gluon distributions, will address key questions in the 2007 Nuclear Science Long Range Plan (DOE 2007). LQCD calculations are crucial to experimental investigations of the hadron structure of nucleons at the Jefferson Laboratory, Relativistic Heavy Ion Collider-spin and a possible future electron-ion collider. Notably, these calculations, together with experiments, will resolve the origin of spin in the nucleon. These calculations will also delineate between the roles of the spins of the quarks and gluons, and of their orbital angular momentum with a precision that neither experiments nor computation can achieve alone.

The progression toward extreme computing for hadron structure is encapsulated in Figure 40. LQCD will enable precision calculations of key isovector quantities. These include the nucleon axial charge, which impacts the lifetime of the neutron, electromagnetic form factors specifying the spatial distribution of charge and magnetization in the nucleon, moments of quark distributions measured in deep inelastic scattering, and moments of generalized parton distributions, which are a major focus of the experimental program at Jefferson Laboratory. Calculations requiring computational resources approximately one

petaflop-year at the physical pion mass are required for these observables to be extrapolated to infinite volume and the continuum with an accuracy of a few percent.

Separate calculation of neutron and proton form factors, moments of quark distributions, and of generalized parton distributions (GPDs) require the calculation of more computationally demanding disconnected quark contributions, originating from the sea quarks. Using algorithms that have recently been developed, 10-100 petaflop-years will enable calculation of these disconnected diagrams and therefore meet the DOE 2014 milestone HP9 (DOE 2008):

> Perform lattice calculations in full QCD of nucleon form factors, low moments of nucleon structure function and low moments of generalized parton distributions including flavor and spin dependence.

This will include detailed imaging of the two-dimensional transverse spatial structure of the nucleon. An outstanding example of synergy with the experiment is the combination of moments of GPDs calculated with LQCD, and convolutions of GPDs measured at the Jefferson Laboratory and elsewhere, which together will provide a more complete understanding than either effort could separately obtain.

Calculations requiring computational resources of order 100 petaflop-years will increase the precision of the axial charge calculated with LQCD to a level of better than 1%, which will begin to impact the calculation of the proton-proton fusion rate central to solar models.

Extreme computation is required for the calculation of nucleon form factors to sufficiently high-momentum transfer to explore the onset of asymptotic scaling behavior. Such calculations will complement the analogous investigations and calculations of structure in, and transition form factors to, unstable baryons such as the $\Delta(1232)$. The calculation of photon structure functions, hadronic polarizabilities, and the exploration of higher moments of structure functions also requires extreme computing resources.

FROM QUANTUM CHROMODYNAMICS TO NUCLEI

In low-energy and low-temperature systems (e.g., conditions as they are on the earth), QCD displays itself through the existence of hadrons (e.g., protons and pions) and their interactions. Exactly how this occurs has been a long-standing question in fundamental physics.

The coupling of effective field theories (EFTs) and LQCD (e.g., Beane et al. 2008c) in recent years has allowed for substantial progress in deriving the interactions between hadrons directly from QCD, particularly for systems involving mesons (e.g., pions and kaons). Similar EFTs, when constrained empirically, have been successful in nuclear many-body calculations of light nuclei (e.g., alpha and lithium). Lattice-based effective field theories (LEFTs) using nucleon degrees of freedom, as opposed to those of quarks, have made large strides in calculating neutron matter, an infinite medium consisting of neutrons, albeit at small densities compared to normal nuclear-matter densities. These research developments are at a nascent stage because of computational limitations. Extreme computing is required to bring research in these areas to full maturity and become a cohesive program.

For example, with extreme computing resources, LQCD calculations of the interactions within multibaryon systems, like the triton and alpha system, will allow for precise extraction of the three-nucleon interaction, a quantity that is currently poorly constrained empirically. EFTs will be constrained

not only empirically, but also by first-principles LQCD calculations, and will be directly fed into many-body nuclear calculations of nuclei, such as those being performed using the no-core shell model and coupled-cluster formalisms. These same theories will be used to calculate neutron matter at larger densities, which will help scientists calculate the properties of nuclear matter—for example, in the outer crusts of neutron stars. This priority research direction (PRD) will therefore have impact not only at the microscopic level, but to earthly and astrophysical phenomena as well.

Basic Scientific and Computational Challenges

Signal-to-Noise. Current LQCD calculations are inherently stochastic. That is, these calculations use random sampling techniques (Monte Carlo) to perform high-dimensional integrals that are necessary to describe physical phenomena. Calculations of baryon systems, such as the deuteron, suffer from poor signal-to-noise ratios due to the stochastic nature of these calculations. This impedes extractions of multibaryon interaction parameters. As computational resources increase, signal-to-noise issues will diminish slightly, but only with the development of novel algorithms and computational techniques can the signal-to-noise issue be resolved. A similar issue, commonly known as the "fermion sign problem," is seen in LEFT calculations of neutron matter.

Scaling Multi-Baryon Codes for High-Performance Capability. Existing algorithms for performing multibaryon calculations are not suited for extreme scale computing, and will have to be modified to take advantage of extreme computing capability. However, such modifications will not be straightforward, and novel algorithms for multibaryon systems need to be developed to optimize use of large computational resources.

Development of Finite-Volume EFTs. Continued development of finite volume EFTs needs to occur so that LQCD calculations can be matched onto theories used by other areas of nuclear physics, such as nuclear structure and reactions, and nuclear astrophysics. Such theories will need to be "pionful," allowing for the (perturbative) determination of the light-quark mass-dependence of the interactions and scattering parameters. This will provide the most robust extrapolation methods.

Interfacing with Large-Scale Nuclear Structure Calculations. LQCD calculations of few-nucleon interaction parameters will ultimately be fed into nuclear many-body calculations via the use of EFTs, such as those being performed with no-core shell model and coupled-cluster theories. EFTs matched to LQCD calculations (typically in a plane-wave basis) will need to be adapted for nuclear structure calculations (which typically use the harmonic oscillator basis). At the two-particle level, this matching is straightforward in that there is an exact analog of Luscher's formalism (Luscher 1986*)* within a harmonic oscillator basis. However, for three- and higher-body nucleon systems, significant research remains to be done. This will entail substantial collaboration with the nuclear structure community—something that is currently just beginning. Both theoretical and numerical methods need to be developed to enhance the overlap between LQCD and the nuclear structure and reactions community. Similar efforts need to be made with the nuclear astrophysics and "hot and dense" QCD scientific communities.

Scientific Outcomes and Impacts

Extreme scale computing will greatly extend the computational prowess of the nuclear physics community. Certain calculations, only aspirations before, will now be accessible and have

transformational impact on the broader physics community as a whole. The following are some key scientific outcomes in this PRD that will result from extreme scale computing resources.

Three-Body Interaction Between Baryons. LQCD calculations of three- and four-baryon systems, such as the triton and alpha particle, will allow for the extraction of various three-body interaction parameters that are currently poorly constrained (if at all) empirically. Of particular importance is the three-nucleon interaction, which has implications to the nuclear structure and reactions community. The three-body interactions between nucleons and hyperons will also be accessible, which in turn could have direct implications in astrophysical settings.

Binding Energy of Alpha Particle. For the first time, the four-nucleon system will be calculated directly from QCD. This particle represents the heaviest s-shell nucleus—its inclusion into the suite of LQCD calculations will allow for a comprehensive constraint on the interaction parameters in the EFTs needed for nuclear many-body calculations. Probing this system will also give scientists insight into the four-nucleon interaction—something that presently cannot be done experimentally.

The impact of the above outcomes, and the research leading up to these outcomes, are widespread. For example, insufficient knowledge of the three-nucleon interaction is responsible for the largest systematic uncertainties in nuclear structure and reaction calculations of light nuclei. Without a better knowledge of this interaction, absolute binding energies and level orderings of excited states of nuclei cannot be calculated with high fidelity. LQCD calculations at the extreme scale will remove this obstacle. Furthermore, research leading to these calculations will help scientists understand the interactions between two- and three-body systems that are not accessible experimentally, but believed to play prominent roles in astrophysical settings (e.g., the interaction between kaons and nucleons). LEFT calculations of neutron matter will be performed at a precision where scientists can quantitatively state the properties of the crust of neutron stars, which are the remnants of Type-II supernovae.

Figure 41 shows anticipated key highlights obtained with high performance computing as extreme computing capability is approached. Baryon-baryon interaction parameters will be computed in the limit of exact isospin symmetry with high precision with sustained petascale resources. With an order of magnitude increase in computational resources, the deuteron axial-charge will be accessible. This is one of the key ingredients constraining certain fusion reactions—within the sun, for example. Finally, at the extreme scale, the three-nucleon interaction can be calculated, as well as the alpha-particle system. Throughout this period, development on EFTs and their overlap with other subfields of nuclear physics will be performed.

The authors of this report emphasize the research impact in this direction will come from calculating observables that are currently inaccessible by experiment, and have great relevance not only to the QCD community, but to the broader nuclear physics community such as those of nuclear structure and nuclear astrophysics.

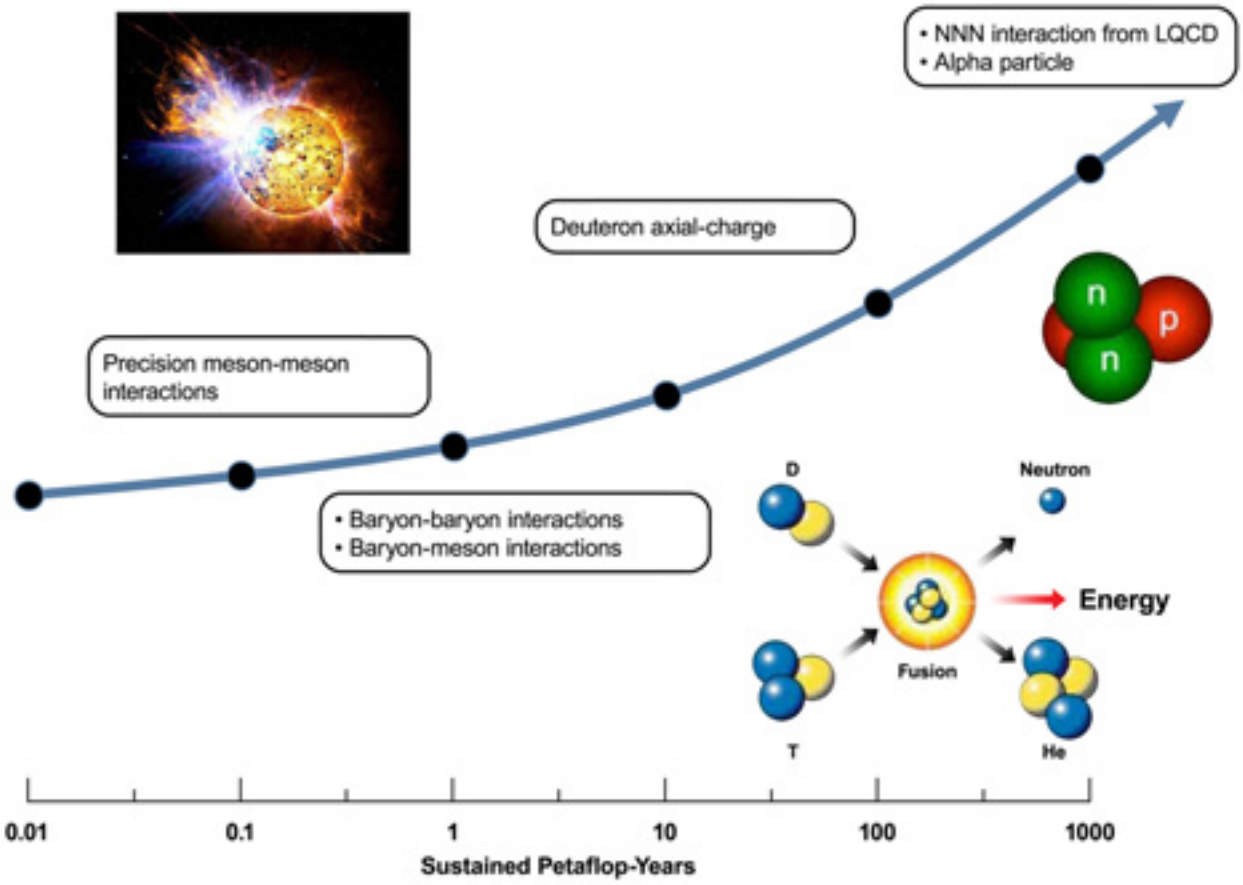

Figure 41. Anticipated highlights for priority research direction "From Quantum Chromodynamics to Nuclei." Upper-left image from NASA. Lower-right image courtesy of the Plasma Physics Laboratory of the Royal Military Academy, EURATOM Association, Belgium. Remaining images courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

FUNDAMENTAL SYMMETRIES

In some instances, nature is very nearly invariant under certain symmetry transformations, such as spatial inversion or motion reversal (also known as time-reversal). However, the consequences of a slight noninvariance under such transformations can have widespread implications. A well known example is CP-violation, where the combined symmetry operation of charge-conjugation, C, and spatial-inversion, P, is known to be slightly violated. Without CP-violation, the present-day matter and antimatter asymmetry of the universe would not exist (the universe contains more matter than antimatter), and from what ensues, humans would not exist.

Research efforts to uncover particles and symmetries beyond those of the standard model of particle physics are multipronged. One of the approaches in this effort is to perform precision measurements of the properties of known particles, such as the magnetic moment of the muon. The E821 experiment at Brookhaven National Laboratory has measured the deviation from the classical value of the muon magnetic moment, g-2, to eight significant digits, and is found to agree with the theoretical calculation within the uncertainties of both the theoretical calculation and the experimental determination. One of the significant uncertainties in the theoretical calculation arises from strong interaction contributions through

quantum loops. Exploratory LQCD calculations are underway to understand the methodology that may be employed to directly calculate these loop contributions.

This PRD aims to quantify the connection between certain violations of fundamental symmetries and the resulting observed physics phenomena. Explicit examples include the parity-violating part of the nuclear interaction, and the electric dipole moment (EDM) of the neutron due to time-reversal violation. This will connect to the DOE 2020 milestone F115 (DOE 2008):

> Obtain initial results from an experiment to extend the limit on the electric dipole moment of the neutron by two orders of magnitude.

Extreme computing will allow, for the first time, a quantitative understanding of how these broken symmetries manifest themselves in nuclear physics interactions. Scientists will gain a much deeper understanding of how these symmetries, at the fundamental level and in particular through electroweak interactions and interactions beyond the standard model of particle physics, impact nuclear physics.

Basic Scientific and Computational Challenges

Four-Point Functions. Calculations of these symmetry-violating observables will generally require a new class of algorithms that enable the calculations of four-point functions. The demands of these calculations will require high-petascale capability, with full maturity coming from extreme scale resources.

Disconnected Diagrams. To date, LQCD calculations have generally involved a certain class of calculations—those of connected diagrams—due to computational limitations. These diagrams represent the propagation of quarks from the initial to the final states. For a complete description of parity violation in nuclear structure, for example, there are short-distance parity-violating few-nucleon forces that contain disconnected diagrams.

Sampling Relevant Topological Sectors. Because symmetry-violations are typically small, the Monte Carlo calculations of these phenomena have signal-to-noise ratios that diminish in time much faster than in most standard LQCD calculations. These "topological fluctuation" issues cannot be remedied by simple reweighting techniques, and will require significant resources to test and develop techniques to avoid this problem.

Memory Requirements. Lattice measurements of all parity-violating effects will have substantial memory requirements due to the large number of distinct light-quark propagators that will be required. Such requirements are currently estimated to be at least two orders of magnitude greater than current available resources.

Scientific Outcomes and Impacts

Parity-Violating Nuclear Interactions. It is generally agreed by scientists that there should be a nucleon-nucleon interaction mediated by one-pion exchange that arises from the weak interaction, and thus a "long-distance" parity-violating contribution to the nuclear force. This parity-violating effect, which is encoded in the weak analogue of the nucleon axial coupling, remains poorly determined despite decades of experimental effort. A LQCD determination of this coupling will have a great deal of impact. In principle, all parity-violating effects in the two-nucleon sector can be calculated with LQCD by

extracting parity-violating two-nucleon scattering parameters from the energy levels of two nucleons in a finite-volume lattice, which interact through four-quark operators. Knowledge of the microscopic origins of parity violation in nuclear physics will help correlate and explain the parity-violating signatures observed in nuclear structure. Moreover, knowledge of the parity-violating nuclear interaction calculated with LQCD will provide an explanation of how the weak interaction at the quark level and the strong interaction conspire to generate weak interaction forces among nucleons, and parity violation in nuclear structure.

Neutron EDM Due to the θ-Term and Higher-Dimension Operators. It is possible that QCD contains CP-violating effects that propagate into the hadronic sector via the so-called θ-term (and also through higher-dimension, "irrelevant" operators). One approach to isolating and quantifying these effects is a direct LQCD measurement of the neutron EDM with a nonzero value of θ. While there have been some preliminary studies, this is a computationally challenging endeavor as the CP-violating effect is expected to be small and its signal quickly diminishes in time as topological fluctuations become smaller in the approach to the chiral limit. An understanding of the presence of CP violation due to the θ-term in QCD will sharpen the search for CP-violation whose origin is beyond the standard model, and more generally, constrains models of physics beyond the standard model.

Figure 42 shows the anticipated milestones of this PRD as the extreme scale computing era is approached. Preliminary calculations of the hadronic parity-violating part of the nuclear interaction will be obtained with petaflop-year sustained resources, followed by the first calculations of the neutron EDM with an order of magnitude increase in computer resources. At the extreme scale, the full effects of the nuclear parity-violating component of the nuclear interaction, as well as the full understanding of the neutron EDM due to the θ-term and higher-dimension operators, will come to fruition. These results will coincide with measurements obtained with experiments at the Spallation Neutron Source at Oak Ridge National Laboratory.

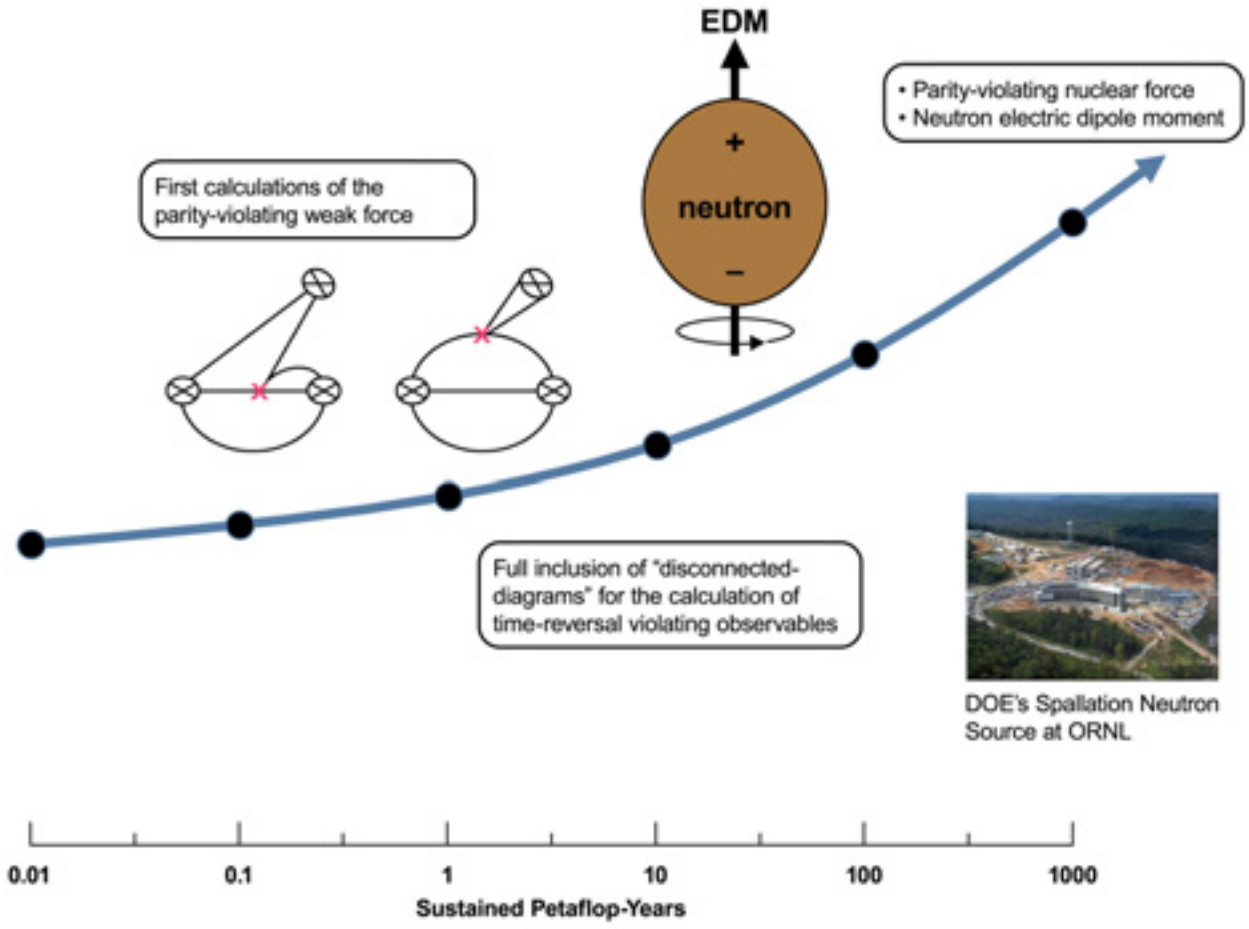

Figure 42. Anticipated highlights for priority research direction "Fundamental Symmetries." Lower-right image courtesy of U.S. Department of Energy, Oak Ridge National Laboratory. Remainder of image courtesy of Thomas Luu (Lawrence Livermore National Laboratory) and David Richards (Thomas Jefferson National Accelerator Facility).

NUCLEAR STRUCTURE AND NUCLEAR REACTIONS

AB INITIO CALCULATIONS OF LIGHT NUCLEI AND THEIR REACTIONS

Basic Scientific and Computational Challenges

A realistic *ab initio* approach to light nuclei with predictive power must have the capability to describe bound states, unbound resonances, and scattering states within a unified framework. Over the past decade, significant progress has been made in understanding the bound states of light nuclei starting from realistic nucleon-nucleon (NN) plus three-nucleon (NNN) interactions (Pieper and Wiringa 2001; Navrátil et al. 2000, 2007; Hagen et al. 2007b). The solution of the nuclear many-body problem is even more complex when scattering or nuclear reactions are considered. For few-nucleon systems (A=2-4), accurate methods solve the bound state and the scattering problems. However, *ab initio* calculations for scattering processes involving more than four nucleons are still the exception (Nollett et al. 2007; Quaglioni and Navrátil 2008; Hagen et al. 2007a) rather than the rule. The development of an *ab initio* theory of low-energy nuclear reactions on light nuclei is key to further refining scientists' understanding of the fundamental interactions between the constituent nucleons. At the same time, such a theory is required to make accurate predictions of nuclear astrophysics' crucial reaction rates that are difficult or even impossible to measure experimentally. This section highlights a key direction that *ab initio* methods will pursue with exascale resources.

REACTIONS THAT MADE US: TRIPLE-ALPHA PROCESS AND $^{12}C(\alpha,\gamma)^{16}O$

Extreme scale computing will enable the first precise calculation of $2\alpha(\alpha, \gamma)^{12}C$ and $^{12}C(\alpha,\gamma)^{16}O$ rates for stellar burning (see Figure 43); these reactions are critical building blocks to life, and their importance is highlighted by the fact that a quantitative understanding of them is a 2010 U.S. Department of Energy milestone (DOE 2007). The thermonuclear reaction rates of alpha-capture on ^{8}Be (2α-resonance) and ^{12}C during the stellar helium burning (see Figure 43 for a schematic depiction) determine the carbon-to-oxygen ratio with broad consequences for the production of all elements made in subsequent burning stages of carbon, neon, oxygen, and silicon. These rates also determine the sizes of the iron cores formed in Type II supernovae (Brown et al. 2001; Woosley et al. 2002), and thus the ultimate fate of the collapsed remnant into either a neutron star or a black hole. Therefore, the ability to accurately model stellar evolution and nucleosynthesis is highly dependent on a detailed knowledge of these two reactions, which is currently far from sufficient.

Experimental measurement of these reaction rates at energies relevant for astrophysics (at approximately 300 keV in the center of mass) is impossible with existing techniques because of their extremely small cross-sections. Because of the influence of alpha-cluster resonances in ^{12}C and ^{16}O, theoretical extrapolations of measurements performed at higher energies to the relevant low-energy region have large uncertainties (for recent measurements, see Assunção et al. 2006). Presently, all realistic theoretical models fail to describe the alpha-cluster states, and no fundamental theory of these reactions exists. Yet, a fundamental theory is needed to determine the rate of the $^{12}C(\alpha,\gamma)^{16}O$ reaction to at least 10% accuracy to fix the subsequent burning stages.

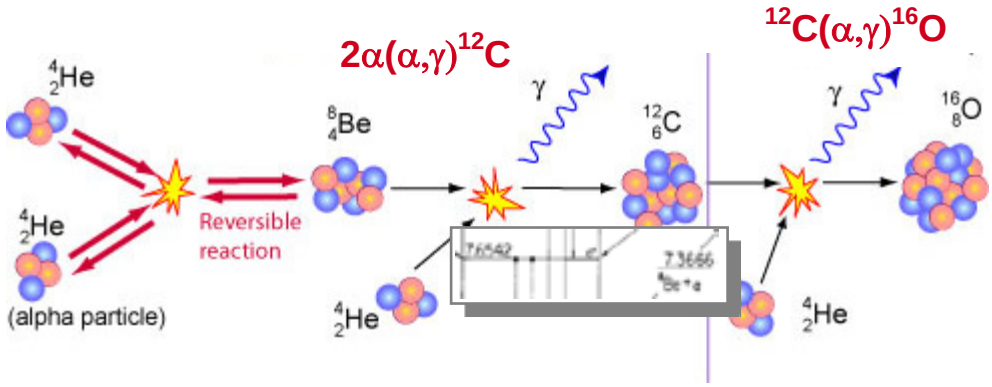

Figure 43. A schematic view of the ^{12}C and ^{16}O production by alpha burning. The $^{8}Be+\alpha$ reaction proceeds dominantly through the 7.65 MeV triple-alpha resonance in ^{12}C (the Hoyle state). Both sub- and above-threshold ^{16}O resonances play a role in the $^{12}C(\alpha,\gamma)^{16}O$ capture reaction. Image courtesy of Commonwealth Scientific and Industrial Research Organisation (CSIRO).

These calculations can be performed by using several independent *ab initio* methods, which will permit results verification and allow for systematic uncertainties to be determined. The methods are as follows: 1) the Green's Function Monte Carlo (GFMC) approach generalized for scattering; 2) the *ab initio* no-core shell model (NCSM) extended by the resonating group method (RGM); and 3) the coupled cluster method. The calculations can proceed in several phases with increasing complexity, and a general picture of the computational requirements for these calculations is shown in Figure 44.

The first phase focuses on the Hoyle state in ^{12}C. This is an alpha-cluster-dominated, 0^+ excited state lying just above the $^{8}Be+\alpha$ threshold and is responsible for the dramatic speedup in the ^{12}C production rate. The calculation of this state will be the first exact description of an alpha-cluster state. It can be achieved with 10% accuracy of the excitation energy within 3 years using the current petaflop machines, and with 5% accuracy in 10 years using improved Hamiltonians. Calculations of alpha-capture on ^{8}Be will be performed within the next 5 years. Calculations for ^{16}O, and in particular of the alpha-cluster resonances that impact the $^{12}C(\alpha,\gamma)^{16}O$ reaction, will follow. Finally, the $^{12}C(\alpha,\gamma)^{16}O$ calculations will be completed within a 10-year time frame.

Scientists can reliably estimate the increase in computer resources needed to address the ^{16}O nucleus with GFMC. Presently, the GFMC calculation of the ^{12}C ground state requires approximately 400 peta operations. The Hoyle state will require tens of calculations of the same size. The number of operations will increase by a factor of approximately 1200 for ^{16}O, with the growth provided by the available computing resources increasing from the petascale to the extreme scale.

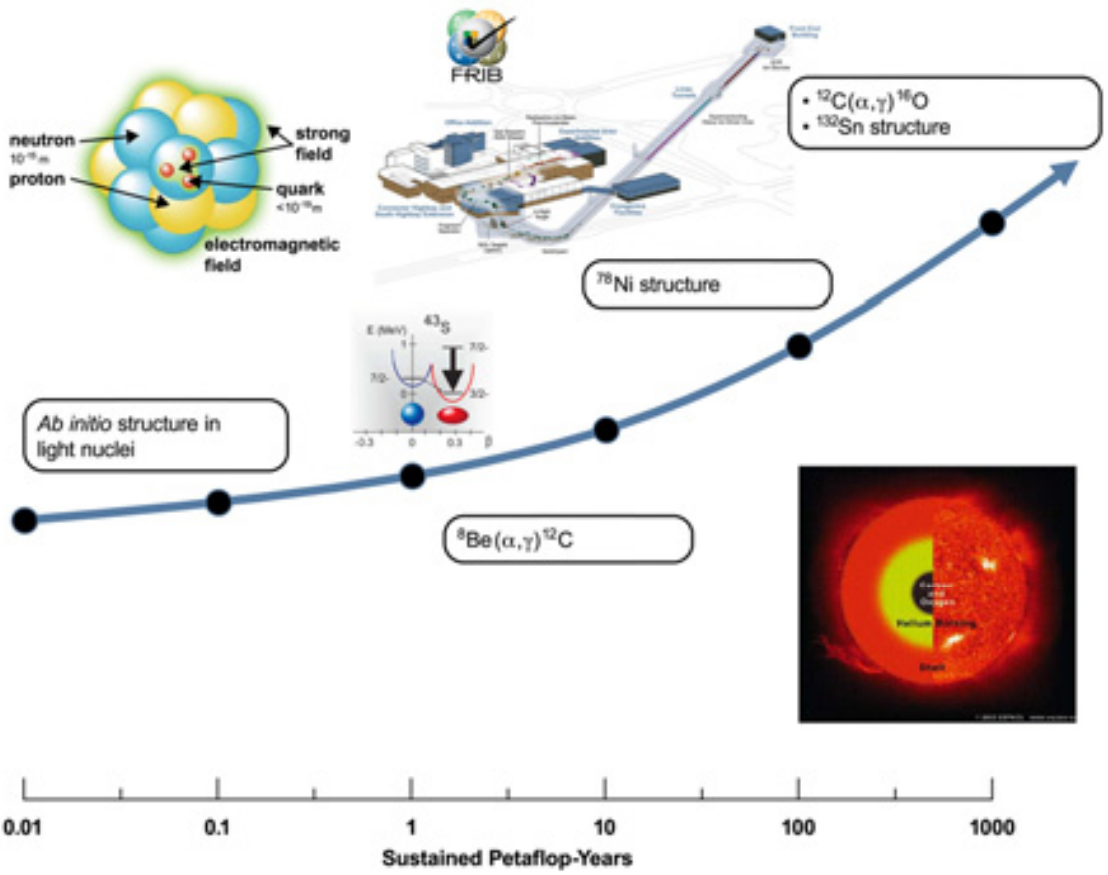

Figure 44. Anticipated highlights for the priority research direction "Reactions That Made Us." Top-middle image courtesy of Michigan State University. Remainder of image courtesy of James P. Vary (Iowa State University).

Currently, it is becoming feasible to calculate, within the NCSM/RGM framework, low-energy nucleon-^{12}C or nucleon-^{16}O scattering with soft NN forces using approximately 1000 cores on present-day machines. The computational demand increases dramatically (a factor of approximately 10^6) with increasing size of the projectile (from a single nucleon to an alpha particle) and by including the NNN interaction. Therefore, this is clearly a problem requiring the extreme scale computation level.

The ground state of ^{16}O can presently be computed within the coupled-cluster 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 than this because of the proximity of the scattering continuum; it will be based 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 not be easy to use an extreme scale computer for these calculations. The present ability in GFMC was obtained by splitting the work on one Monte Carlo configuration among tens of cores (previously just one core was used). For ^{16}O, 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.

In the NCSM/RGM, the matrix elements for hundreds of density operators must be calculated. These calculations are both central processing unit and memory intensive. Calculations are presently completed using message-passing interface (MPI) with distribution of the memory allocation. For example, in the calculation of matrix elements of the density operators, the cores are divided into groups, each of which is responsible for computing matrix elements of a subset of operators. This type of parallelization will need to be optimized and propagated to a finer level of distribution among clusters of computing cores in extreme scale machines as the complexity of the task grows rapidly with the mass of the target nucleus, the mass of the projectile (alpha particle), and presence of the NNN force.

Scientific Outcomes and Impacts

The primary outcome of this effort will be a comprehensive understanding of the mechanism behind these two key reactions, and the ability to model the chemical evolution of the universe. Success will permit an accurate determination of the reaction rates at low energies relevant to stellar burning, which are currently limited by large experimental uncertainties. In particular, the uncertainty in the $^{12}C(\alpha,\gamma)^{16}O$ reaction rate is currently about 40% (Angulo et al. 1999). By achieving this research goal, scientists will enhance the predictive power of stellar modeling. At the same time, scientists will develop *ab initio* tools to describe the structure of weakly bound nuclei that will be studied at the Facility for Rare Isotope Beams (FRIB) at Michigan State University. Verification of model predictions by experiments at FRIB will provide necessary checks on the theoretical approaches and the underlying two- and three-body forces used. One computational outcome will be the development of a library for distributing shared-memory work to subsets of nodes within a massively parallel machine.

A successful completion of this program will provide essential input for modeling of stellar evolution and element production. It will provide a firm basis for extrapolating future experimental results. It will guide and be validated by light exotic nuclei studies at FRIB and other exotic beams facilities. Finally, scientists will understand how ^{12}C and ^{16}O, elements critical for life, are produced in nature.

WEAK NUCLEAR STRUCTURE—NUCLEI AS LABORATORIES FOR NEUTRINO PHYSICS

Basic Scientific and Computational Challenges

The neutrino is one of the most elusive particles in the universe and yet one of the most influential. The mass and interaction of the neutrino with other matter are less than a millionth of an electron's—yet neutrinos power spectacular core-collapse supernovae that seed the universe with heavy elements. The fact that neutrinos even *have* a mass is one of the great discoveries of the past 10 years. If the neutrino is its own antiparticle—a so-called Majorana particle—then physics beyond the current standard model of elementary particles must be invoked with consequences impinging upon the matter and antimatter imbalance in the early universe. While the $^{12}C(\alpha,\gamma)^{16}O$ reactions reveal how life can exist, a Majorana neutrino may reveal how matter itself came to exist. However, stringent upper-limits on the existence of neutrino Majorana mass contributions would force scientists to look to other explanations for the fundamental matter-antimatter asymmetry that is observed in the universe today. For recent reviews of current knowledge of neutrino properties, see Avignone et al. (2008), Camilleri et al. (2008), and Haxton (2008).

The primary venue for discerning the fundamental properties of neutrinos is atomic nuclei. A number of experiments are being planned worldwide to determine their properties, but interpreting the results of those experiments will require reliable calculations of nuclear structure and of the interaction between neutrinos and nuclei. Two broad classes of experiments are relevant here, and because of the difficulty in obtaining constraints needed to calibrate these experiments, both require sophisticated theory to be interpreted. As a check on the calculations, as well as a determination of the systematic uncertainty in the theory, scientists will use competing methods to compute the reaction and decay rates.

The first method consists of long-baseline experiments to measure neutrino flavor oscillations, which are sensitive to the differences in neutrino masses, as well as neutrino flavor-mixing angles. Detectors used in these experiments are based on target nuclei such as carbon and oxygen, and it is crucial to understand the neutrino-induced response of these nuclei to fully exploit measurements. At lower energies, neutrino cross-sections on these nuclei also play an important role in late-stage stellar evolution, as well as driving gravitational-collapse supernovae and the creation of heavy elements in supernovae. Reliable calculations require accurate treatment of the strong interaction and a realistic representation of the weak interaction currents. At low energies, the neutrinos couple with nuclei predominantly through so-called "allowed" operators, which are simple, easily calibrated, and cross-checked through experimentation. However, at higher energies—including those relevant to the detectors—scientists also need "forbidden" current operators, which are much more difficult to compare directly to the experiment.

The second experimental methodology consists of neutrinoless double-beta decay ($0\nu\ \beta\beta$ decay) measurements (Vogel 2007). These decays can only occur if the neutrino is its own antiparticle; if so, a neutrino can be emitted and reabsorbed within the same nucleus. If these decays do occur, the lifetime is inversely proportional to the mass of the neutrino and the nuclear matrix element. Unlike $2\nu\ \beta\beta$ decay, which can be largely calibrated by comparison to ordinary beta decay, the operator responsible for the $0\nu\ \beta\beta$-decay nuclear matrix element is neither theoretically simple nor easily constrained by other experiments. Among the specific target nuclei are ^{48}Ca, ^{76}Ge, and ^{130}Te.

The fundamental challenge is to create a computer model of the structure of a nucleus, and then compute the nuclear coupling to neutrinos. Starting from fundamental measurements of NN interactions and using rigorous mathematical methods, effective interactions suitable for use on petascale and extreme scale computers, as well as the weak current operators that describe the interactions of neutrinos with nucleons, will be developed. This will be a significant computational project. A general illustration of the computational requirements for these calculations is provided in Figure 45.

Two main techniques will need to be extended to use extreme scale computing facilities: quantum Monte Carlo (QMC), primarily for the ν-nucleus cross-sections and configuration-interaction shell model (CI-SM) for the $0\nu\ \beta\beta$-decay nuclear matrix element and subsequent lifetime. These have complementary strengths and weaknesses. QMC techniques can use "bare" NN and NNN interactions taken directly from experiment, but QMC cannot yet tackle the heavy nuclei, such as ^{76}Ge or ^{130}Te, relevant to $0\nu\ \beta\beta$ decay. CI-SM is the technique of choice for detailed spectra and can use arbitrary forms of interactions, not just local potentials—but to fit the problem even on an extreme scale machine, the NN interaction must be renormalized. A third technique, quasi-particle random-phase approximation, makes computationally much more modest demands and is thus widely used, but is a more severe approximation.

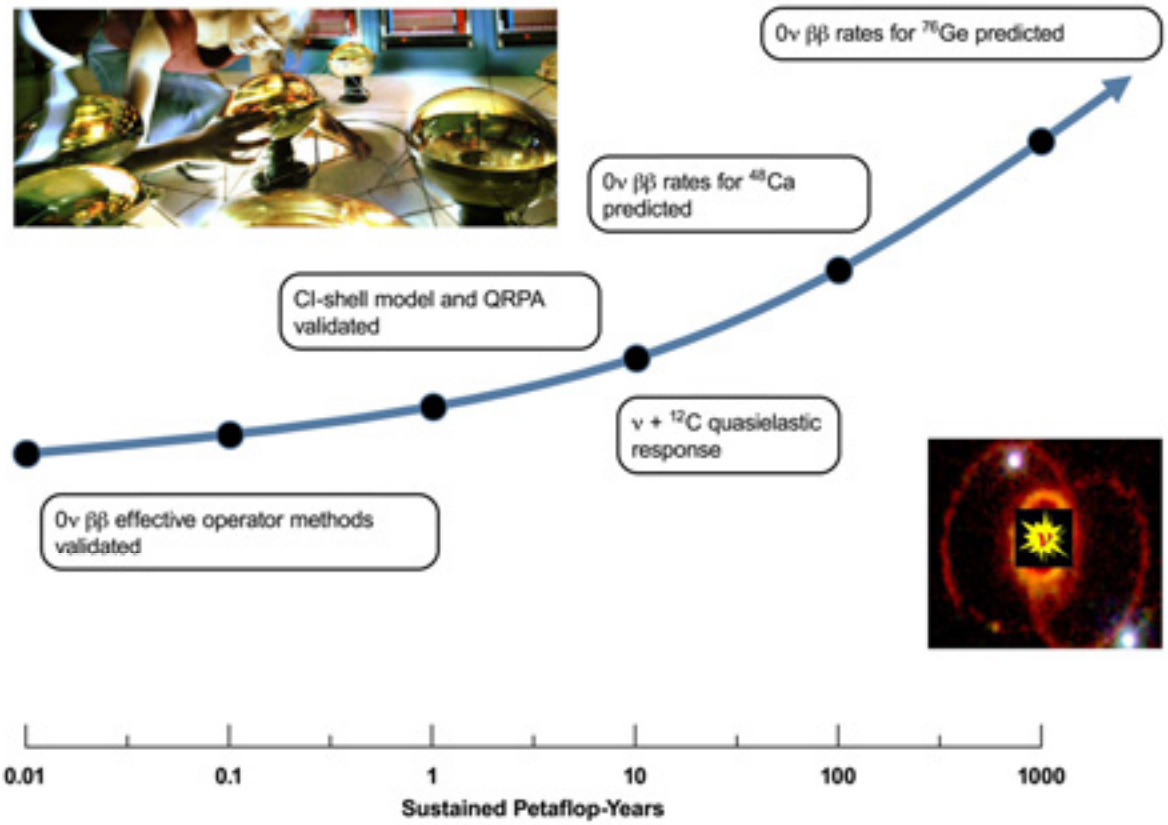

Figure 45. Anticipated highlights for priority research direction "Nuclei as Neutrino Physics Laboratories." Image courtesy of James P. Vary (Iowa State University).

Each of these techniques faces challenges to be scaled to extreme scale computers. QMC techniques must have actions that are now confined to a single computing core distributed over multiple computing cores. CI-SM will require finding the lowest part of the spectrum of a very large matrix, with dimensions on the order of 1-10 trillion; although the matrix is very sparse, storing the nonzero elements will require petabytes of memory. Furthermore, CI-SM requires vector operations that must communicate across the entire machine. Finally, for CI-SM, scientists must renormalize the experimentally determined NN interaction; this in itself will be a computationally intensive problem because one needs to evaluate the induced NNN interactions in large-basis spaces.

Scientific Outcomes and Impacts

If the neutrino is its own antiparticle, the resulting 0ν ββ-decay lifetime of various nuclei will depend sensitively on the absolute mass of the neutrino. The goal is to compute the 0ν ββ-decay lifetime for nuclei relevant to planned experiments with theoretical uncertainty to 30-50%, cross-checked using competing methods. Accurate estimates of the expected lifetime could affect design of experiments requiring expensive isotopically enriched materials. If a 0ν ββ-decay lifetime is actually measured, these calculations will enable the extraction of the neutrino mass.

Long-baseline oscillation experiments measure the difference between neutrino masses as well as other parameters of the neutrino mass matrix. To correctly interpret the experiments, the ν-nucleus cross-sections will be required to be computed with uncertainties that are less than approximately 20%.

CI-SM can also compute neutrino cross-sections that can provide a cross-check of the QMC and quasiparticle random-phase approximation calculations. As part of this computational project, calculations will be compared using several different methods, usually with the same starting point, from which a systematic uncertainty associated with the calculation can be estimated. One important issue for CI-SM is renormalization, not only of the interaction between nucleons but also between neutrinos and nucleons. Rigorous renormalization methods exist and must be applied consistently to the interaction and the neutrino coupling. Comparisons with results from QMC, where more direct models of the current can be employed, will provide crucial validations.

Currently, significant experimental effort and funds are being invested to answer the above questions, but the experimental results cannot be persuasively evaluated without significant theoretical effort. With extreme scale computing, theoretical studies will provide a basis for reliable interpretation of experiments that explore the properties of neutrinos.

MICROSCOPIC DESCRIPTION OF NUCLEAR FISSION

Basic Scientific and Computational Challenges

Current understanding of nuclear fission, a fundamental nuclear decay, is still incomplete because of the complexity of the process. Nuclear fission has many societal applications ranging from power generation to national security. In addition, it also plays a role in the synthesis of nuclei in the r-process. Yet, to date, scientists have no microscopic understanding of this complex phenomenon and are unable to make reliable and accurate predictions of fission half-lives, cross-sections, or the distribution of fission products. The ongoing (2009) Scientific Discovery Through Advanced Computing Program (SciDAC)-2 Universal Nuclear Energy Density Functional project (Bertsch et al. 2007) and petascale computing resources are opening the way for a comprehensive microscopic description of static properties of atomic nuclei and the fission process.

A promising starting point to obtain a predictive model of nuclear fission is the density functional theory (DFT); see the Nuclear Fission Extreme Scale Computing sidebar. This theory provides the justification for an energy-functional approach to explaining and predicting nuclear structure across the complete table of the nuclides. The accurate nuclear energy functionals currently in use are purely phenomenological and have parameters that are fit to only a subset of nuclear properties. Petascale computing resources and improvements in DFT codes made available through the Universal Nuclear Energy Density Functional project (Bertsch et al. 2007) are opening avenues to the comprehensive microscopic description of complex nuclear phenomena in general, particularly in nuclear fission. Several approaches, each entailing a number of serious computational challenges, can be applied to the description of nuclear fission and will be pursued in this program. The adiabatic approach requires as a first step the determination of the potential energy surface (PES) in a multidimensional space of collective coordinates, which comes from constrained Hartree-Fock-Bogoliubov (HFB) calculations (Warda et al. 2002; Staszczak et al. 2005).

Including all relevant degrees of freedom to obtain a realistic and precise PES is a particularly challenging task. Compounding this issue is the need to evaluate the inertia tensor (Giannoni and Quentin 1980;

Warda et al. 2002; Goutte et al. 2005). For this program to succeed, it will be critical to develop suitable algorithms to improve the efficiency of constrained calculations. The imaginary-time HFB (Levit et al. 1980; Arve et al. 1987; Puddu and Negele 1987; Skalski 2008) approach relies on the computation of the full spectrum of dense complex matrices with dimensions that can reach millions. Not all of these matrices are Hermitian. Solving eigenvalue problems of that scale will require an enormous amount of memory, which will create a major bottleneck in the calculations. Heterogeneity in future computer architectures (e.g., use of graphical processing units) will pose another complication. New approaches will therefore be needed to overcome the memory bottleneck in these extreme scale calculations. A general illustration of the computational requirements for these calculations is provided in Figure 46.

Fission half-lives are extremely sensitive to the details of the underlying PES and the collective mass tensor. This requires extending the current program of energy density functional development to an unprecedented level of precision because phenomenological energy functionals provide essentially a qualitative description. Novel functionals will typically involve 10-30 parameters to be determined through the global minimization of a large number of observables. Constraining effectively each term of the energy functional requires performing symmetry-unrestricted HFB calculations and possibly adopting techniques beyond the mean-field methods. The dimensionality of the problem, combined with the necessity to reach the global minimum, will probably require massive global optimization algorithms. The phenomenon of fission will be investigated with various microscopic approaches. A first step from current capabilities is to follow the adiabatic time-dependent Hartree-Fock-Bogoliubov theory.

At least four degrees of freedom—elongation, mass asymmetry, necking, and triaxiality—must be considered. To attain sufficient mesh refinement, it will be necessary to compute of the order of 100,000-plus constrained HFB calculations for every nucleus.

Two nonadiabatic approaches will also be explored. The first is the instanton method, which relies on determining periodic trajectories for the imaginary time HFB equations. Finding the bounce solutions (periodic instantons) is a difficult numerical challenge. The second approach, applicable in the context of induced fission where the explicit time propagation can be conducted, is a stochastic extension of the time-dependent superfluid local density approximation (TD-SLDA) of DFT. The appeal of this approach, equivalent to the many-body Schrödinger equation, is that two-body and higher correlations become accessible, and dissipation is naturally incorporated into the theoretical description. TD-SLDA has been successfully implemented on current leadership-class super computers, specifically on the Cray XT4 Jaguar at Oak Ridge National Laboratory. A stochastic realization of TD-SLDA will require sufficiently large ensembles of size from thousands to millions of realizations.

Nonadiabatic approaches to spontaneous and induced fission will allow the prediction of the mass and excitation energy distribution of the fission fragments, half-lives, and cross-sections. Beyond the scission point, the emerging fragments start accelerating, and the binding energy of the mother nucleus is converted partially into the kinetic energy of the fragments. At the same time, because strong dissipative processes become increasingly more important, a significant part of the energy is converted into the internal excitation energy of the fragments. The stochastic approach to the time-dependent fission dynamics will allow scientists to calculate these dissipative processes microscopically and predict the nuclear viscosity.

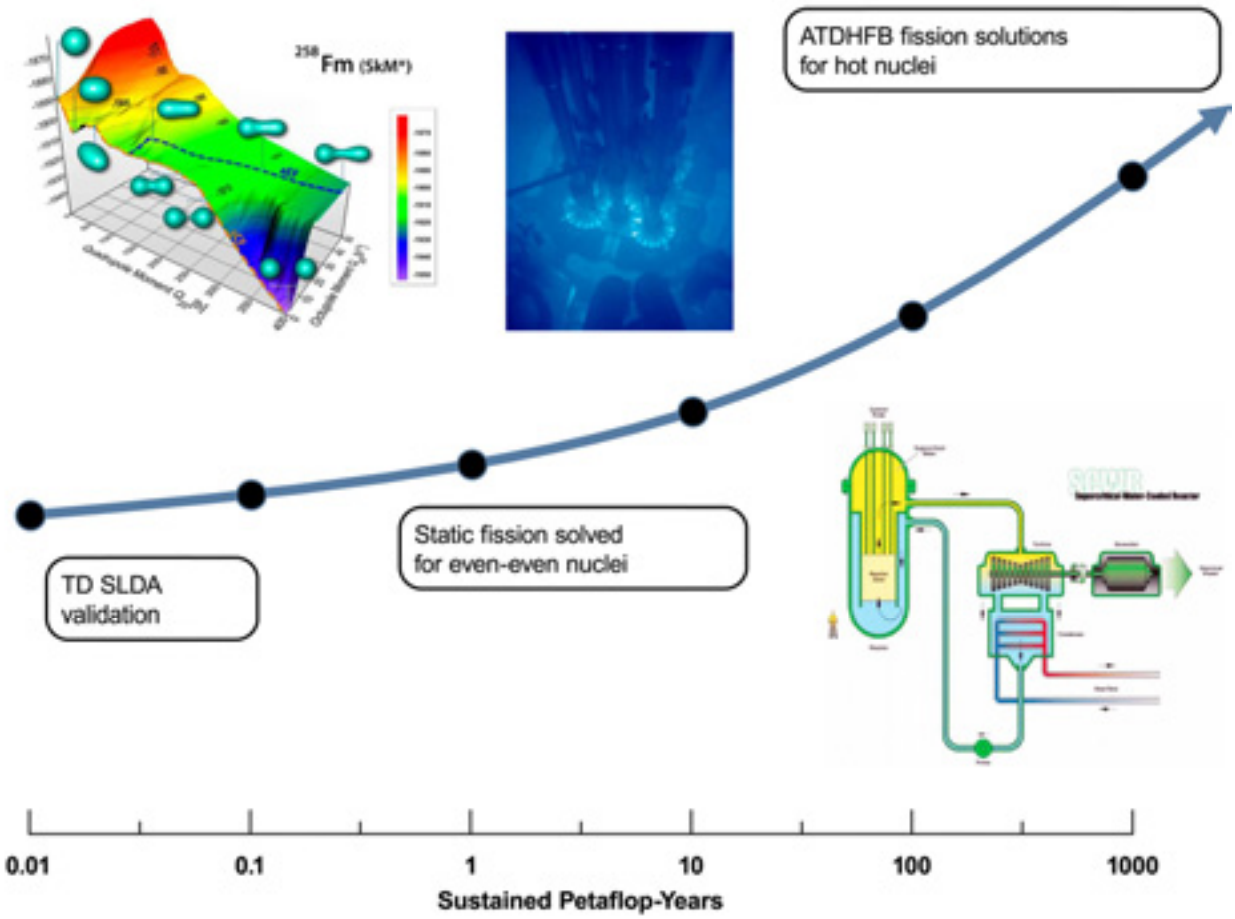

Figure 46. Anticipated highlights for priority research direction "Microscopic Description of Nuclear Fission." Top-left image courtesy of A. Staszczak, A. Baran, J. Dobaczewski, and W. Nazarewicz (Oak Ridge National Laboratory). Remainder of image courtesy of James P. Vary (Iowa State University).

One of the implementation difficulties of stochastic TD-SLDA is the large local memory demand per MPI process and the limited random-access memory/core. Current state-of-the-art calculations prescribe a single MPI process per node so that all the memory in a node is aggregated into a larger, addressable local memory. This approach leaves the other processor cores idle or requires lightweight thread level control within the MPI process to use these cores. Scientists anticipate the need to increase the size of the Hilbert spaces, which will exacerbate this memory-aggregation problem or force the computations out of core—effectively stalling productivity even in the single determinant problems. Programming techniques that go beyond single-node memory aggregation will be refined or developed to satisfy this need. Such developments will also need to include the implicit/explicit use of the extra processor cores.

Scientific Outcomes and Impacts

The computational approach to fission envisioned here, combined with experiments, will provide a predictive framework that may lead to improved nuclear reactor design (AFC 2006). In the area of national security, developing a theoretical description of fission aligns with the goals of the National Nuclear Security Administration Stockpile Stewardship Program, which entails an accurate and complete modeling of the behavior and performance of devices in the nation's aging nuclear weapons stockpile.

Improving the accuracy of that description is central to the continuing process of certifying both the safety and the reliability of the stockpile without resumption of nuclear testing and to reduce the threat from nuclear proliferation.

Of all the various nuclear decay processes, nuclear fission—important in the r-process nucleosynthesis, in the modeling of reactions relevant to the advanced fuel cycle for next generation reactors, and in the context of national security—is among the most difficult to tackle. It is a quantum many-body tunneling problem whose typical time-scale changes by orders of magnitude when adding just a few nucleons. The microscopic theory of nuclear fission, rooted in internucleon interactions, still provides a particularly difficult challenge.

The ultimate outcome of the nuclear fission project is a treatment of many-body dynamics that will have wide impacts in nuclear physics and beyond. The computational framework developed in the context of fission will be applied to the variety of phenomena associated with the large amplitude collective motion in nuclei and nuclear matter, molecules, nanostructures, and solids.

PHYSICS OF EXTREME NEUTRON-RICH NUCLEI AND MATTER

Basic Scientific and Computational Challenges

Understanding neutron-rich nuclei is vital to discovering the origin of heavy elements (NAP 2003) and defining the properties of neutron-star crusts (Ravenhall et al. 1983). About half of the elements from iron to uranium are produced via successive steps consisting of neutron capture followed by beta decay (the r-process). The structure of neutron-rich nuclei determines the radiative capture cross-sections and beta-decay rates that are critical inputs to r-process nucleosynthesis calculations. The regions around the supposed doubly magic nuclei ^{60}Ca, ^{78}Ni, and ^{132}Sn are of particular interest as they could be waiting points in the r-process. The existence and location of shell closures affect the r-process path as illustrated in Figure 47, where the r-process path is schematically drawn assuming shell closures at the traditional magic numbers. The dynamic and static properties of neutron star crusts determine neutron-star cooling and gravity wave emissions from neutron star mergers.

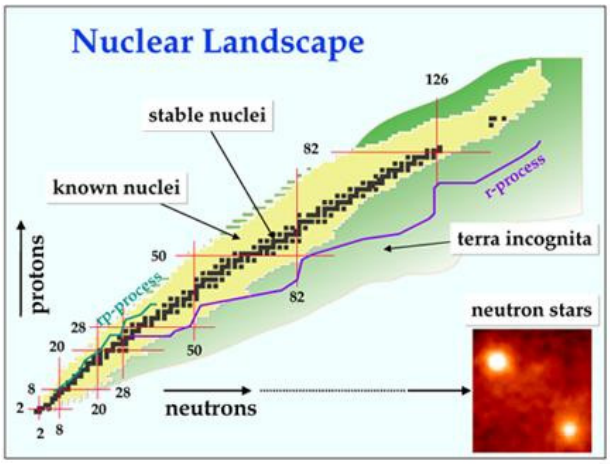

Figure 47. The chart of atomic nuclei displays the speculated r-process path of rapid neutron capture across neutron-rich nuclei. The structure of extremely neutron-rich nuclei is essential input to understand the origin of heavy elements as well as the cooling properties of and the gravity wave emission from neutron star crusts. Image courtesy of James P. Vary (Iowa State University).

Unfortunately, present understanding of neutron-rich nuclei is very limited, and extrapolations based on current theoretical models are not reliable. First, the extreme isospin of neutron-rich nuclei magnifies unconstrained properties of the effective nuclear interaction. Second, the proximity of the neutron drip line dramatically increases the number of relevant many-body configurations, including the continuum, and makes accurate computations impossible at the present time. In the coming decade, progress towards the most neutron-rich nuclei will be

made with both theory and experiment. The future FRIB at Michigan State University will provide experimental data for selected nuclei along the r-process path. These data will calibrate and validate theoretical methods which, with the advent of exascale computing facilities, will enable accurate theoretical predictions for extremely neutron-rich nuclei (see Figure 47).

The *ab initio* nuclear-structure program aims at building nuclei starting with nucleon degrees of freedom and their mutual interactions. Extending this program to neutron-rich nuclei in the ^{60}Ca, ^{78}Ni, and ^{132}Sn regions and towards the neutron drip lines poses great theoretical and computational challenges. A general picture of the computational requirements for these calculations is illustrated in Figure 48.

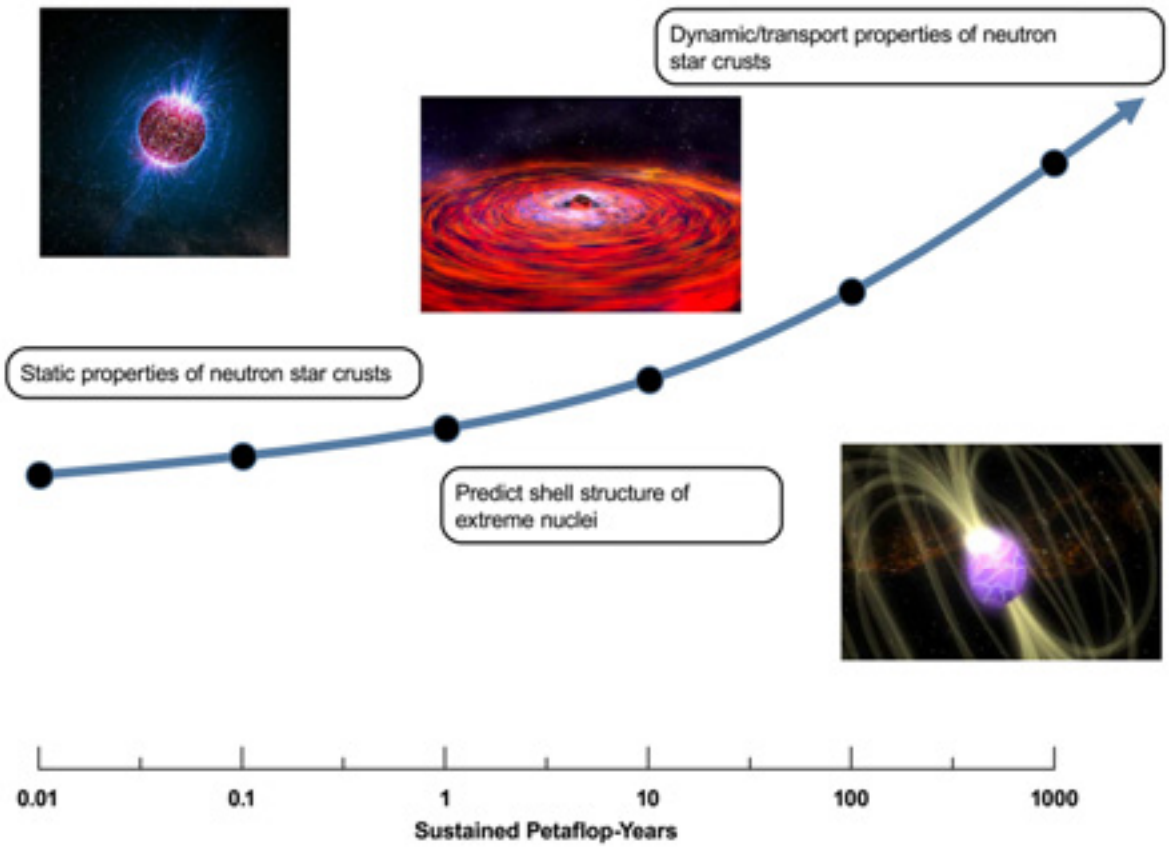

Figure 48. Anticipated highlights for priority research direction "Physics of Extreme Neutron-Rich Nuclei and Matter." Image courtesy of James P. Vary (Iowa State University).

Closed-shell nuclei and their neighbors are of particular interest for both experimental and theoretical research because they form the pillars of understanding and modeling for atomic nuclei.

The effective nuclear Hamiltonian, including the isospin dependence of the effective nuclear two- and many-body forces, is under intense investigation and will become far more precise in the next 3 years. These interactions will be employed with state-of-the-art nuclear-structure tools such as configuration interaction (Lisetskiy et al. 2004), the coupled-cluster method (Hagen et al. 2008), the nuclear density-functional theory (Bertsch et al. 2007), and Monte Carlo techniques (Chang et al. 2004) to calculate the properties of closed-shell nuclei and their neighbors. Of particular interest are the regions around the neutron-rich nuclei ^{78}Ni and ^{132}Sn. These calculations will predict the evolution of shell structure and will

explore the drip line and the limits of nucleonic matter. For the understanding of neutron star crusts, the transport properties of systems composed of extremely neutron-rich nuclei and a surrounding neutron gas must be computed.

Calculations of nuclei in the ^{78}Ni region and of static properties of matter in the crust of a neutron star require a facility with tens of petaflop-years of capacity, while computations of nuclei in the ^{132}Sn region and transport properties of crust matter require a facility with hundreds of petaflop-years capacity. Scientists assume the program will be balanced such that investments in computational hardware and software are matched with investments in theory and personnel.

Scientific Outcomes and Impacts

Calculations of nuclei in the ^{78}Ni region and of static properties of matter in the crust of a neutron star require a facility with tens of petaflop-years of capacity, while computations of nuclei in the ^{132}Sn region and transport properties of crust matter require a facility with hundreds of petaflop-years capacity. Scientists assume the program will be balanced such that investments in computational hardware and software are matched with investments in theory and personnel.

These extreme scale computations will allow scientists to determine the limits of nuclear stability—that is, how many neutrons or protons can be bound in a given nucleus. This theoretical effort will have a major impact upon the experimental program to search for these limits at research facilities such as the FRIB. The combination will allow scientists to model some of the most exotic environments in astrophysics, and understand and model the chemical evolution of the universe.

In the crust of neutron stars, neutron-rich nuclei coexist with a surrounding gas of neutrons; the structure and dynamic properties of this unusual matter will be calculated using advanced Monte Carlo methods. In turn, it will be possible to interpret the wealth of astronomical data obtained from visual, x-ray, and gamma-ray telescopes. This will allow scientists to infer details of the nature of these sites and the processes (such as potentially gravitational wave emission) that occur there. The major computational challenge in these efforts is to develop and implement scalable algorithms for the strongly interacting inhomogeneous quantum many-body problem.

An important complement to the work described here will be the experimental program conducted at the FRIB (NRC 2006). The theoretical and computational tools envisioned above will provide an essential framework to interpret FRIB experimental data and will eventually guide the future experimental program. In turn, FRIB data will be essential to verify *ab initio* calculations and calibrate the nuclear many-body Hamiltonian.

Computations of neutron star matter, when combined with observations, will provide information about nucleonic matter at supernuclear densities. The interpretation of observations of isolated, cooling neutron stars require an accurate microscopic understanding of superfluidity and neutrino emission processes in neutron-rich matter. Similarly, observations of gravity waves with the advanced Laser Interferometer Gravitational-Wave Observatory and future detectors will, when combined with a realistic description of the neutron star matter, allow scientists to infer the mass and radius of a neutron star. Combined observations of multiple neutron stars will produce definitive constraints on the equation of the state of the densest matter in the universe.

NUCLEAR ASTROPHYSICS

SUN AND OTHER STARS

Basic Scientific and Computational Challenges

While estimates of nucleosynthetic yields are widely used to infer the origin of the elements that compose the sun and the solar system, as well as to predict the behavior of abundances in the first stars, they depend strongly on the treatment of hydrodynamic mixing in turbulent regions inside stars. Turbulence is a notoriously challenging phenomenon but is ubiquitous in stellar interiors—therefore, a deeper understanding is essential for developing a predictive theory of stellar evolution. Turbulent mixing is a significant problem during the late stages of evolution (post-carbon burning), at which time the nuclear evolution in the stellar core decouples from the observable surface properties of the star and calibrating the physics of mixing is not possible. Using calibrated mixing rates based on earlier phases of evolution is also not guaranteed to apply during the distinct vigorous core- and shell-burning convection accompanying the late burning stages.

Computationally, modeling stellar turbulence strains the presently available resources because of the enormous range of relevant length and timescales and the variety of physical processes involved. In supernova progenitors and helium-shell flash nucleosynthesis, reactive hydrodynamic flows need to be modeled. This requires a multiple-component fluid description to track the compositional evolution and the associated nuclear energy release, both of which in turn feed back into the dynamics through buoyancy forces. In the case of solar convection, modeling the photosphere involves multiangle, multigroup radiative transfer, which adds considerable extra cost. In addition, magnetic fields are likely to play a nonpassive dynamic role so that a magnetohydrodynamic solution is desirable.

Capturing a high enough Reynolds number (i.e., Re > 1000) is the biggest obstacle that must be overcome to reliably model stellar flows. In a simulation, the effective Reynolds number scales with linear zoning across a domain as $Re \propto N^{4/3}$ for turbulent flow. Therefore, it would be ideal to have approximately 180 zones across each of the relevant scale lengths (large energy-containing eddies) that arise in the flow. In the solar convection zone, the lower convective boundary layer, known as the tachocline, is approximately 10 times narrower than the convection zone depth. Therefore, to resolve this transition layer, approximately $\sim 2{,}000^3$ zones spanning the entire region would be ideal. Fewer zones may be sufficient if an informed choice of nonuniform zoning is used, thereby decreasing the needed zone count by a factor of a few in each dimension. This would lead to an overall reduction of zones by an order of magnitude. Adaptive mesh refinement is not as important as a nonuniform grid for stellar interior modeling because turbulence at high Reynolds numbers is space filling, and scientists are generally interested in studying quasi-equilibrium states. Scalable adaptive mesh refinement methods, however, can provide the underlying computational framework needed to employ a fixed mesh refinement grid on a massively parallel architecture. Additional savings, perhaps as large as $1/\mathcal{M} \approx 100$ in computing time, may be achieved for low-Mach number flows $\mathcal{M} \approx 10^{-2}$, if scalable low-Mach number methods are successfully developed for petascale and exascale platforms.

For each of the three problems highlighted, breakthroughs will be made possible in moving from the petascale to the extreme scale, primarily because of the increased computational volumes and degree of turbulence (i.e., Reynolds number) achievable. In solar convection modeling, petascale resources will afford enough resolution that a turbulent tachocline can be self-consistently incorporated into a global model. Such a simulation would provide a breakthrough in the scientists' ability to understand the heat and angular momentum transport, which is mediated by this boundary layer and determines the differential rotation profile and dynamo action observed in the sun. Extreme scale computational resources would allow for the self-consistent modeling of solar surface granulation within a global circulation model of the sun, providing precision tests of both the simulation techniques and an understanding of the global-scale magnetohydrodynamic activity observed in the active sun.

Developing three-dimensional supernova progenitor models involves a large range in both spatial and temporal scales. While the end state of a massive star depends upon the entire prior evolution of the star since formation, a three-dimensional simulation that begins at core silicon burning and is evolved up to core collapse would provide a significantly improved level of confidence about the state of the iron core at collapse, including the rotational state and the convectively induced perturbations. Such a three-dimensional stellar model would be used directly as an input to core-collapse supernova simulations. Spatially, capturing global asymmetries will require simulating a volume that extends to the outer edge of the carbon burning convection zone (Meakin and Arnett 2006). Thus, in successively larger shells surrounding the core, silicon, oxygen, neon, and carbon burning will need to be included. Angular momentum transport by wave motions in the stable layers between convection zones (e.g., Talon and Charbonnel 2005) are likely to be important during this epoch and will require a similar computational volume for study. At the petascale, a two-dimensional model encompassing the carbon burning shell could be undertaken and would provide a first-generation multidimensional supernova progenitor model. In a three-dimensional model, the properties of the silicon burning core could be simulated for an hour preceding collapse, thus incorporating realistic features of the vigorous convection. With extreme scale computing resources, a three-dimensional model for the entire silicon burning epoch that incorporates all of the overlying burning shells out to carbon burning would be possible.

The timescale relevant to s-process nucleosynthesis in asymptotic giant branch stars is set at a minimum by the period over which helium shell burning convection persists, which is approximately 10 years, while the time period between helium outbursts is approximately 10^3 years. For comparison, the convective turnover time is approximately 3 hours (Herwig et al. 2006) and the Courant time (the hydrodynamic time step limit) is a factor of $f \propto N_{zones}/\mathcal{M} \sim 10^4$ times smaller still, for a Mach number $\mathcal{M} \approx 0.01$ and a modest $N_{zones} \approx 100$ zones spanning the convective shell. Scientists are therefore faced with the problem of evolving the model for an extraordinarily large number of time steps. This temporal problem can be ameliorated by studying snapshots of the quasi-equilibrium turbulent flow. These snapshots guide basic theory to be implemented in stellar evolution codes. This approach requires a three-dimensional simulation spanning only 10 to 100 convective turnovers (Meakin and Arnett 2007). In addition to this temporal challenge is the spatial challenge of capturing the flow in the overlying convective envelope, which extends to very large radii ($r_{env} \approx 3 \times 10^{12}$ cm) compared to the size of the helium burning shell ($r_{shell} \approx 10^9$ cm). While the petascale would allow a first-generation three-dimensional model with sufficient resolution to capture a turbulent convective envelope, the extreme scale would allow for three-dimensional giant star simulations that achieve resolved boundary layer mixing over secular timescales. This is the essential jump from the ability to calculate the bulk to the ability to calculate the surfaces, interfaces, and fine details that yield the observable shape.

The computational developments that would benefit stellar interior modeling include the following: 1) low-Mach number techniques that are scalable to mega-core platforms (e.g., Lin et al. 2006; Almgren et al. 2006), and 2) improvements in reaction network solvers that are informed by reduced quasi-equilibrium-group physics (Hix et al. 2007; Arnett 1996) that can more efficiently treat the complex silicon burning epoch preceding core collapse in massive stars. Treating the solar photosphere self-consistently in a deep convection simulation entails a radiation-hydrodynamics problem, which would benefit from techniques capable of load balancing the multigroup, multiangle radiation transport methods (Nordlund 1982) on a mega-core computing architecture. Finally, a data management challenge is inevitable because of the long integration times necessary to obtain the robust statistics required for studying quasi-steady, turbulent flow. A typical petaflop-scale turbulence simulation with approximately $2{,}000^3$ zones that is sampled 100 times per large eddy turnover for two turnovers would generate approximately 10 petabytes of data if stored at single precision. The total data generated, D, for a turbulence model taking advantage of the available flop rate, F_R, will scale roughly as $D \propto F_R^{3/4}$, so that at the exascale, data volumes should be on the order of a single exabyte.

Scientific Outcomes and Impacts

Stellar evolution, including stellar death through supernovae, answers the question of the origin of the elements in the cosmos. With a firmer knowledge of mixing in stars, the field of stellar evolution theory and observation will be elevated to that of a precision laboratory for studying the systematics of nuclear matter under extreme conditions, including heavy element nucleosynthesis. The observed solar neutrino flux is already providing important constraints on weak interactions and neutrino oscillation parameters. Extreme scale computing platforms offer exciting new prospects to address several outstanding issues in nuclear astrophysics connected to stellar evolution. Three key areas that will greatly benefit include the following: 1) conducting solar hydrodynamics, 2) performing supernova progenitor modeling, and 3) mixing and nucleosynthesis in giant stars. Figure 49 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

The sun plays a special role as a test of stellar evolution theory because its physical parameters are so well measured. Scientists know its mass, age, radius, and luminosity. Helioseismology has mapped the sound speed to an accuracy of better than 0.5% throughout most of the sun. Solar neutrino spectroscopy has determined the solar core temperature to about 1%. A combination of photospheric and meteoritic measurements constrains solar composition. Scientists can observe the sun's magnetic activity and measure its surface emissions and differential rotations. Modeling this star develops an understanding of the environment on earth and gives scientists an important physics laboratory as the solar neutrino story (Davis 2003) so clearly illustrates.

The one-dimensional standard solar model leaves out many phenomena believed to be important to the sun including convective zone activity, the depletion of light elements in the photosphere, mixing near the radiative/convective zone boundary, and the early convective core—a consequence of out-of-equilibrium carbon burning. The deficiencies of this model are becoming more apparent. Recent three-dimensional modeling of the photosphere—which greatly improved the general consistency of absorption line analyses—has led to reductions in key abundances.

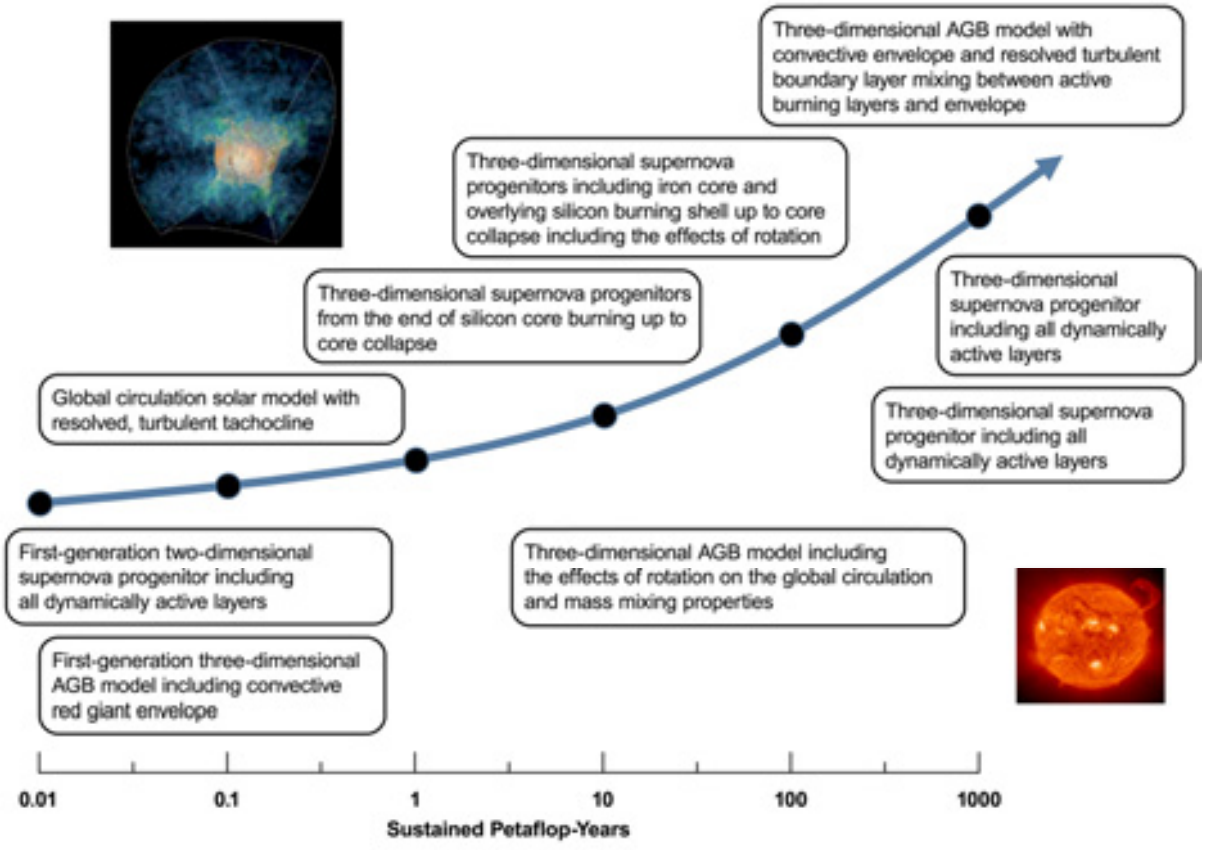

Figure 49. Anticipated highlights for priority research direction "The Sun and Other Stars." Image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

In the standard solar model, these abundances must be used throughout the sun, leading to significant changes in sound speeds and a conflict with helioseismology. The differences are most dramatic in the upper radiative zone, where there could be convective overshoot mixing to alter the structure. Alternatively, the standard solar model assumption of homogeneity at zero age main sequence might be incorrect. Ideally, this assumption could be replaced with an explicit three-dimensional calculation of proto-solar formation through collapse of the primordial solar system gas cloud. One speculation connects the photospheric abundance problem with late-stage formation of the planets, which swept out massive quantities of metal from the nebular disk. Therefore, the time is appropriate to bring the level of realism attainable with extreme scale computing to this important astrophysics problem. Aspects of the standard solar model that could be altered in three-dimensional models, including the initial distribution of core metals and the rate of heavy element diffusion, could alter the fluxes of certain neutrino species by up to 20%, limiting the accuracy of the extractable fundamental neutrino parameters. Supernova progenitors evolved in three-dimensional models, which will be possible at the extreme scale, will serve as initial data for both core-collapse and thermonuclear supernova simulations and will provide a significantly improved level of realism over the one-dimensional models that are currently being used. These simulations will provide insight into the complex interplay between convection and weak interactions (the Urca process), which has consequences for the thermal state of a stellar core prior to explosion. The simulations will also address issues related to symmetry breaking by convection. This symmetry breaking in turn seeds instabilities during the supernova event and informs scientists of the

rotational state of the stellar core prior to core collapse in massive stars. These tests scenarios proposed to explain gamma-ray burst explosions, which require rapidly rotating cores (MacFadyen and Woosley 1999). In the case of asymptotic giant branch stars, placing knowledge of the mixing occurring in the burning shells and envelopes of these stars on a more solid base enables the production of a predictive model of heavy element formation. These predictive models, used in concert with the copious observational data of the surface abundances in these stars, will provide a powerful laboratory for better understanding element synthesis in the cosmos.

The impact on basic nuclear data is far reaching, as the field of stellar evolution has led to several comprehensive compilations of nuclear data that are used widely in the astronomical community. These data compilations represent active fields of research and have led to standards in the field—such as the rate libraries of Rauscher and Thielemann (2000), which provide a means for assembling experimental and theoretical nuclear physics data from a widely dispersed global effort. These standard libraries of data are easily accessible and enable astrophysicists to explore the broader implications of developments in nuclear theory. For instance, experimentally measured nuclear properties—such as neutron-separation energies and neutron-capture Q values (e.g., see Baruah et al. 2008), and experimentally measured reaction rates, such as $^{12}C(\alpha,\gamma)^{16}O$ (e.g., Assunção et al. 2006),[1] have far-reaching consequences for stellar evolution models and nucleosynthesis (Weaver and Woosley 1993; Tur et al. 2007). Astronomical observations of the abundance patterns across the cosmos—including those in the sun and solar system materials, such as meteorites; low metallicity stars; and giant envelopes—make contact with this input nuclear physics data and are interpreted explicitly through the scenarios outlined by stellar evolution theory. In addition, observations of specific signatures of nuclear physics, such as the presence of radioactive nuclides in giant envelopes (Cameron 1955; Gallino et al. 1998) and the spectrum of solar neutrinos (Bahcall et al. 2001), are examples of the direct contact that can be made between stellar theory and nuclear physics.

STELLAR EXPLOSIONS AND THEIR REMNANTS: THERMONUCLEAR SUPERNOVAE

Basic Scientific and Computational Challenges

Explosions of Type Ia supervovae (SNe Ia) involve hundreds of nuclei and thousands of nuclear reactions. These explosions also involve complex hydrodynamic phenomena taking place in degenerate matter and strong gravitational fields (rendering terrestrial experiments of limited utility). Buoyancy-driven turbulent nuclear combustion during the deflagration phase dominates the early part of the explosion and drives an expansion and pulsation of the star. A deflagration-to-detonation transition (DDT) and propagation of the resulting detonation wave through the star has been posited to explain the observed nucleosynthesis and its distribution in space and velocity (Nomoto et al. 1984; Khokhlov 1991; Gamezo et al. 2005). In the alternative gravitationally confined detonation model, fluid flow triggers a detonation that sweeps through the star, producing the observed abundances, spatial distribution, and velocities of the elements.

All of this takes place in approximately 3 s, followed by rapid free expansion of the star at velocities of 10,000 - 25,000 km s^{-1}. These phenomena involve spatial scales from approximately 10^{-3} cm - 10^{9} cm

[1] See panel report titled, "Nuclear Structure and Nuclear Reactions," in this report for further detail on this key reaction.

and temporal scales from approximately 10^{-10} s - 10 s, making simulations of SNe Ia a manifestly exascale problem. Advances are needed in both the speed at which the problem can be addressed and the scale (physical size) of the system that can be handled.

Several key physical processes in SNe Ia are not fully understood, and consequently the understanding of the explosion mechanism is uncertain. These physical processes include the smoldering phase, which precedes the explosion phase and is thought to determine the number of points where ignitions occur and their location(s). The buoyancy-driven, turbulent nuclear combustion phase—or deflagration phase, which releases nuclear energy and expands the star—also represents a frontier, as the understanding of reactive turbulence in strong gravity is incomplete. Finally, the origin of the detonation wave that incinerates the star and causes it to explode is uncertain. Whether the physical conditions necessary for a DDT are achieved in the deflagration phase of SNe Ia is unclear. The alternative, in which fluid flow during the deflagration phase triggers the detonation, is not fully understood.

Extreme scale computing resources will produce breakthroughs in the understanding of these physical processes, transforming scientists' ability to simulate SNe Ia. It will enable the qualitative improvement of scientists' understanding of the smoldering phase, thereby reducing the uncertainty in the initial conditions for simulations of the explosion phase. It will make possible studies of buoyancy-driven turbulent nuclear combustions—which include capturing this physical process by simulations that resolve length scales only three to four decades below the largest physical scales and that use a self-similar subgrid model (Khokhlov 1995; Zhang et al. 2007) if needed—that could verify current expectations. If these studies do not verify these expectations, extreme scale computing will determine that the process is more complicated and provide the data needed to construct an appropriate subgrid model. Finally, extreme scale computing will also make possible studies that verify whether buoyancy-driven turbulent nuclear burning in a white dwarf star produces the physical conditions needed for a DDT to occur.

Advances in SNe Ia modeling during the next decade will most likely come from a combination of high-resolution simulations of the key physical processes described above and whole-star simulations of SNe Ia. Sustained petascale computing will enable verification studies of buoyancy-driven turbulent nuclear burning that will dramatically improve the understanding of this key physical process and will make possible whole-star SNe Ia simulations to treat buoyancy-driven turbulent nuclear combustion over a larger range of scales, providing new insights into the energy cascade and instabilities produced by this physical process. With extreme scale computing, it may be possible to attempt first-principle simulations of SNe Ia from ignition through the deflagration phase (i.e., the buoyancy-driven turbulent nuclear burning phase), a difficult problem.

A key component of studies at both the petascale and the extreme scale will be global validation of the models using large numbers of SNe Ia simulations. The need to perform large ensembles of simulations means the average time to perform simulations of adequate resolution will have to be reasonably short to allow for several such simulations to be performed in a given real time. Thus, a careful mix of a few high-fidelity and many low-fidelity simulations will be required. Even so, it means that high-capacity as well as high-capability extreme scale computing platforms will be needed.

Achieving extreme scale computing capabilities for SNe Ia simulations presents several challenges:

- SNe Ia simulation codes need to exhibit strong scaling and run efficiently on platforms with millions of cores and/or that exploit accelerators. Weak scaling will be insufficient because the computational demand scales as the fourth power of the resolution.

- SNe Ia simulations, in common with core-collapse supernova and stellar evolution simulations, require many physical variables per computational cell (e.g., fluid variables, flame variables, nuclear species variables, and radiation transport variables). Thus, the smaller memory per core of future platforms will require the development of new algorithms for efficient domain decomposition and load balancing.
- New parallel input/output algorithms need to be developed. These include those that can handle files of many terabytes and beyond, along with mass stores that can accommodate exabytes of data. The turbulent nature of the deflagration phase demands high-temporal resolution in the retained data sets. This leads to the production of remarkable data volumes (many petabytes and possibly exabytes).
- New algorithms for scientific data analysis, including visualization, need to be developed to process petabytes and exabytes of data, along with data archiving techniques that can process exabytes of data and allow for comparative analyses to be performed between huge data sets.

Scientific Outcomes and Impacts

The major scientific outcomes of SNe Ia simulations at the extreme scale will be as follows: 1) ascertaining the explosion mechanism; 2) calibrating SNe Ia as standard candles to an accuracy sufficient to study quantitatively the behavior of dark energy with redshift (i.e., with the age of the universe); and 3) understanding the contribution of SNe Ia to nucleosynthesis. Figure 50 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

Understanding the explosion mechanism will also impact ideas about the interaction of reactive flow and turbulence. The deflagration phase is ultimately a straightforward problem in combustion, trading many of the complications of terrestrial burning (e.g., geometry of devices, unmixed fuels, soot production, etc.) for far more fundamental ones (e.g., extremely strong gravity, huge Reynolds numbers, and remarkably stiff reaction kinetics). As such, SNe Ia simulations represent unique numerical laboratories in which to explore basic ideas in reactive turbulent flow. The production of realistic SNe Ia simulations will require advances in this basic area.

An understanding of the explosion mechanism will make possible simulations that can predict correlations among the observed properties of SNe Ia. This will allow them to be better calibrated as standard candles, enabling them to be used to study quantitatively the behavior of dark energy with redshift, and thus to have a strong impact on scientists' understanding of dark energy.

SNe Ia simulations also predict the nucleosynthetic yields for various elements and isotopes, yields that can be tested by observations. These yields are intimately connected with the physical processes that occur during the explosion phase. Consequently, comparisons of nucleosynthetic predictions with observations provide indirect information on these processes, and therefore on the explosion mechanism.

With carbon and oxygen burning being followed by silicon burning and, in the deep interior, an extended period in nuclear statistical equilibrium, SNe Ia simulations are voracious consumers of the nuclear data that govern these burning processes, including binding energies; partition functions; and strong, electromagnetic, and weak interaction reaction rates (e.g., Calder et al. 2007; Seitenzahl et al. 2009). Important reactions, like $^{12}C(\alpha,\gamma)^{16}O$ and triple-alpha burning to form ^{12}C, are the target of ongoing efforts to better measure their reaction rates.

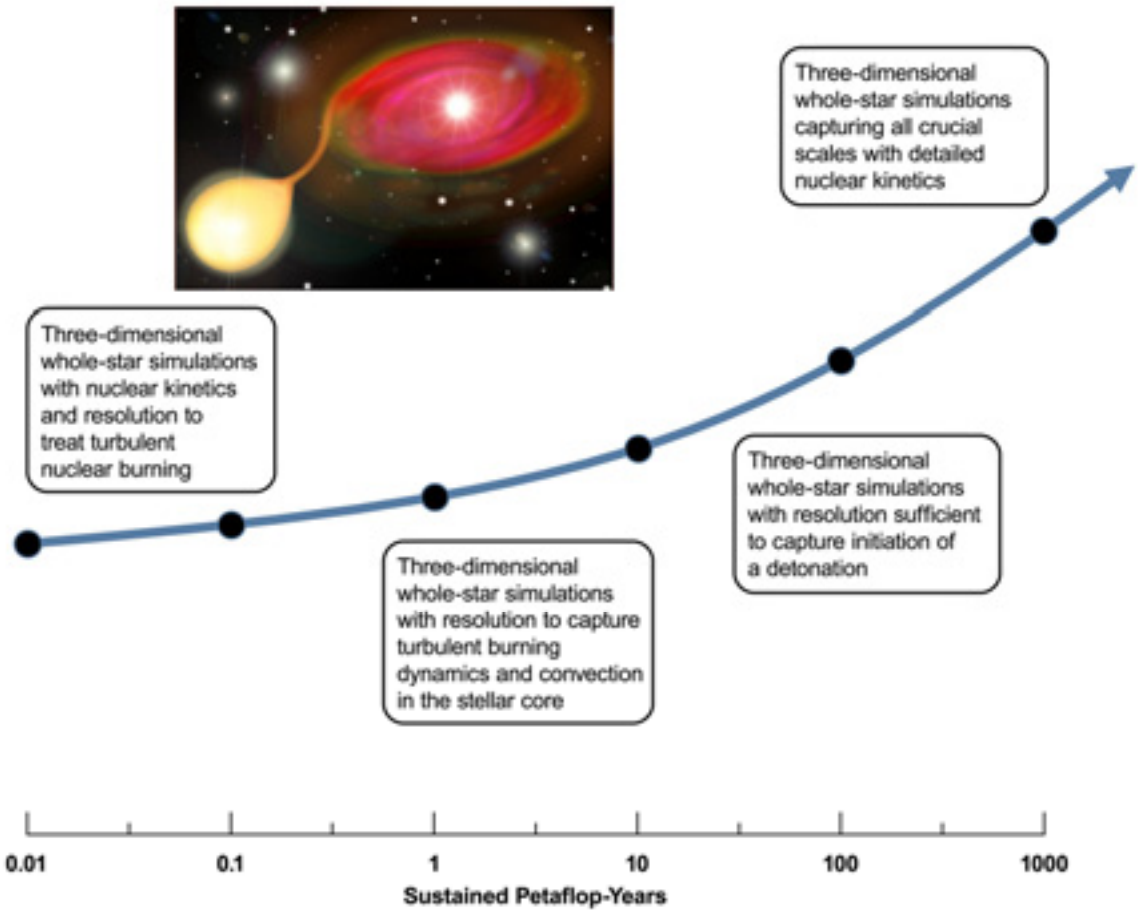

Figure 50. Anticipated highlights for priority research direction "Stellar Explosions and their Remnants: Thermonuclear Supernovae." Upper-left image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

Of particular importance are the weak interaction rates for isotopes of iron peak elements, which determine the neutron richness of the simulated ejecta. These continued improvements in the nuclear data improve the nucleosynthetic predictions from SNe Ia simulations, thereby strengthening the constraints on them that are imposed by observations of their ejecta and of solar abundances.

STELLAR EXPLOSIONS AND THEIR REMNANTS: CORE-COLLAPSE SUPERNOVAE

Basic Scientific and Computational Challenges

Iron-core collapse and bounce are governed by the interplay of general relativistic gravity with the weak and strong nuclear interactions at extremes of neutron richness and density (e.g., $> 10^{14}$ g/cm^3). The subsequent evolution of the event involves neutrino radiation hydrodynamics and nuclear kinetics among other physical processes. The experimental fact of nonzero neutrino masses means scientists must ultimately solve a macroscopic-scale problem in quantum kinetics as well, directly computing the

dynamic, nucleosynthetic, and other observational consequences of flavor oscillations *in situ* as part of fully integrated simulations.

There are profound consequences to this complexity. These simulations make use of a variety of computational algorithms and implementations and stress essentially all facets of a modern, general-purpose computer—the input/output, memory size and latency, processor performance, communication bandwidth and latency, and more—in a manner shared with only a handful of other computational problems. These simulations will stress all facets of a general-purpose supercomputer.

The ability to simulate core-collapse supernovae realistically will depend on the development of discrete representations of the underlying nonlinear partial differential and integro-partial differential equations governing their evolution. This will require efficient and scalable-solution algorithms of the resultant nonlinear algebraic equations, as well as computer codes based on these solution algorithms that can take advantage of the memory and central processing unit capabilities of petascale to extreme scale architectures. Advances in each of these areas will be required, along with considerable work devoted to enhancements of the computational ecosystem surrounding these machines. Core-collapse supernova codes produce prodigious volumes of simulation data over long periods of time. Efficiently writing these data and managing and analyzing them after they are written are as important to producing meaningful science through supernova simulation as is any algorithmic or implementation improvement that might be made for the computational step itself.

Using current petascale platforms and their immediate successors, scientists may be able to determine the general nature of the explosion mechanism itself by performing three-dimensional radiation-magnetohydrodynamics simulations with spectral neutrino transport. As machines capable of peak speeds of 100 petaflops emerge, significant quantitative statements concerning the details of explosive nucleosynthesis in the event and the neutrino emission can be expected. At the extreme scale, scientists will finally be able to determine precisely how supernovae explode by undertaking transformative numerical experiments that incorporate quantum kinetics on macroscopic scales with nuclear physics components realistic enough to accurately predict the isotopic output of these events. These kinds of simulations are utterly unimaginable on current platforms but promise to be accessible at the extreme scale. This is truly applying quantum mechanics, a theory of the smallest things known, to some of the most “macroscopic” bodies in the universe.

Scientific Outcomes and Impacts

The multiphysics nature of core-collapse simulations will require new computational techniques ranging from scalable linear algebra to methods to solve coupled ordinary differential equations. The high number of degrees of freedom at each spatial grid point (e.g., neutrino flavors, energies, and angles, as well as nuclear species) currently represents a large amount of unrealized parallelism in modern supernova codes. Methods to handle these calculations concurrently on multicore platforms and platforms incorporating accelerators of various kinds will likely determine the efficacy of future codes. Figure 51 shows the anticipated key research highlights obtained with high-performance computing as the extreme computing era is approached.

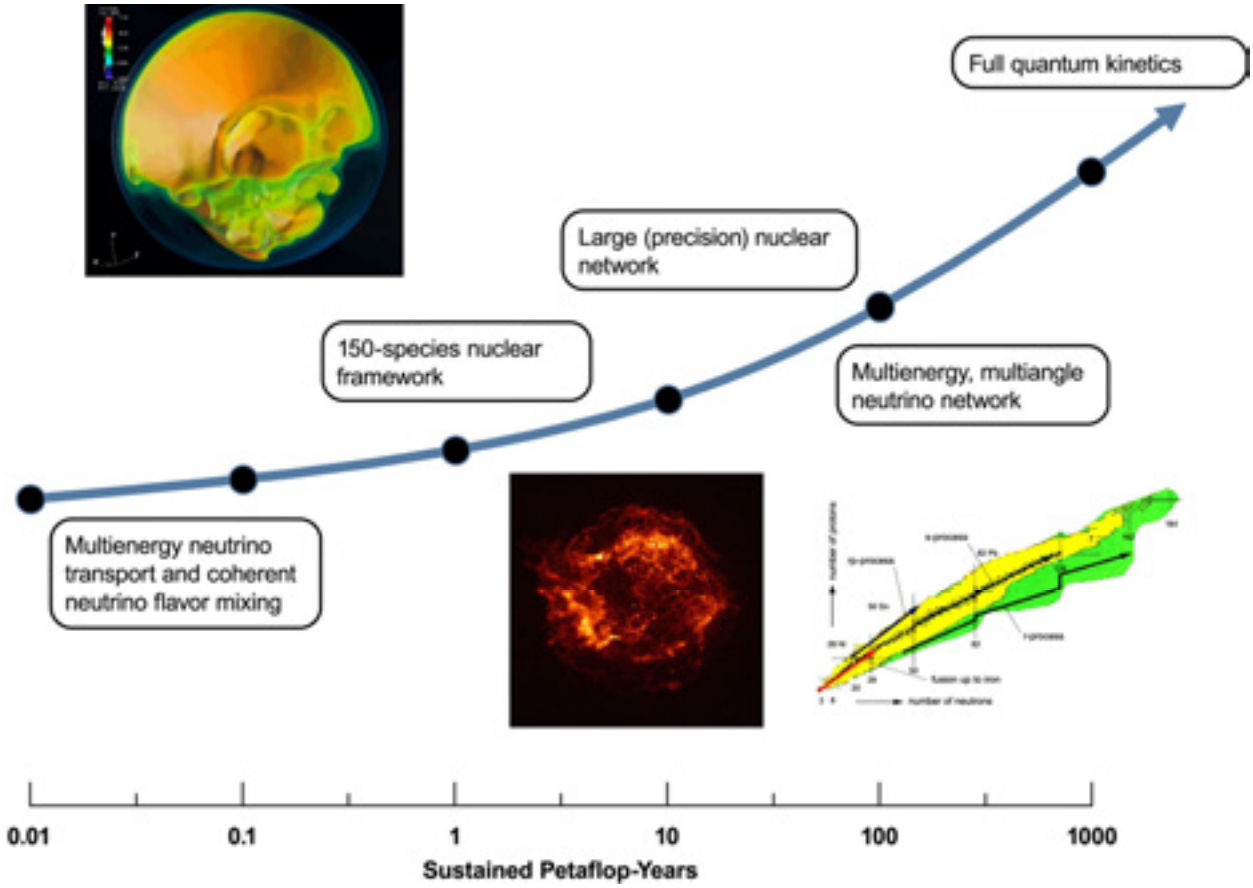

Figure 51. Anticipated highlights for priority research direction "Stellar Explosions and their Remnants: Core-Collapse Supernovae." Bottom-left image courtesy of Chandra X-Ray Observatory and NASA. Remainder of image courtesy of Anthony Mezzacappa and Bronson Messer (Oak Ridge National Laboratory) and George Fuller (University of California).

Several of the major questions posed in the 2007 Nuclear Science Long Range Plan (DOE 2007) are germane to core-collapse supernova simulation.

What are the phases of strongly interacting matter, and what roles do they play in the cosmos?

What is the nature of neutron stars and dense nuclear matter?

The nature of dense nuclear matter formed at the center of a supernova explosion provides a unique opportunity to explore the low-temperature, high-density region of the quantum chromodynamics phase diagram. Knowledge obtained from observation and simulation in this region will complement the better-studied, high-temperature (e.g., quark-gluon plasma) regions of the phase diagram, which are presently accessible to terrestrial experiment.

What is the origin of the elements in the cosmos?

One of the most important and distinctive observables from core-collapse supernovae is their pattern of nucleosynthesis. The creation and transmutation of a wide variety of intermediate- and high-mass species in the event is a nonlinear phenomenon. Supernova nucleosynthesis has a dynamic effect on the

explosion mechanism, ultimately rendering post-processing of simulation results to be of only qualified utility. The subsequent dissemination of the produced species enriches the interstellar medium, setting the stage for successive generations of star formation and death.

Nuclear physics experiments at the Facility for Rare Isotope Beams, combined with improvements in nuclear theory, will constrain temperature, density, timescales, and neutrino fluxes at the r-process nucleosynthesis site from observations of elemental abundances (RIA Working Group 2006). Simulations of core-collapse supernovae will be the essential ingredients in connecting these experimental measurements to the astrophysical site of the r-process, because a self-consistent determination of all of these conditions can only be achieved through computation at scales beyond those currently possible.

What is the nature of the neutrinos, what are their masses, and how have they shaped the evolution of the universe?

Core-collapse supernovae are, from an energetics point of view, neutrino events. They represent the only instance in the modern universe where neutrino interactions have a discernible, macroscopic effect on the dynamics of baryonic matter. Spectral neutrino transport is required to accurately model the event, and the resulting neutrino templates will be invaluable in interpreting and calibrating detections in terrestrial experiments. Comparing future observations to simulation results will be vital to interpreting those observations and using them to constrain the properties of neutrinos.

Accurate and precise knowledge of the characteristics of neutron-rich matter at high density is a prerequisite for understanding core-collapse supernovae. Precise data for electron-capture processes on progressively larger nuclei is a fundamental need for the simulations, a need which can only be filled by advances in nuclear structure theory. Conversely, core-collapse supernova simulations provide the crucial link in testing these theoretical results, as it is only at the extremes of density and neutron richness realized in these simulations where these predictions are manifest. As nuclei in the collapsing core make the transition from an ensemble of nuclei to nuclear matter, exotic forms of matter are expected (Ravenhall et al. 1983). The details of this transition region are of considerable importance in determining accurate neutrino spectra, again providing a unique link between fundamental theory and physical observables.

In addition, core-collapse supernovae are prodigious sources of gravitational waves (GW) (Ott 2009). Because the signal-to-noise ratio for GW detectors presents a serious complication for detection, the production of useful templates for detectors, such as the Laser Interferometer Gravitational-Wave Observatory and VIRGO, is essential for meaningful data analysis. Furthermore, as nonaxisymmetric oscillations are required for GW production, multidimensional, fully integrated simulations are the only path to producing these signal templates. Therefore, the only path forward to interpreting possible future GW wave detections from core-collapse supernovae relies wholly on simulations providing the requisite context.

HOT AND DENSE QUANTUM CHROMODYNAMICS

PRECISION CALCULATION OF BULK THERMODYNAMICS

Basic Science and Computational Challenges

Establishing the properties of matter in the vicinity of the chiral phase transition, and characterizing their dependences upon the quark masses and the number of quark flavors will provide fundamental insight into the many remarkable features of quantum chromodynamics (QCD). It will enable a study of the interplay between the confinement of quarks and gluons and asymptotic freedom, and a study of the role played by chiral symmetry breaking and topological excitations in generating the masses of the hadrons. Furthermore, establishing the properties of strongly interacting matter in the limit of zero net baryon number density is a prerequisite for any further analysis of the QCD-phase diagram at nonvanishing net baryon number density.

To have complete theoretical control of the thermodynamics of strongly interacting matter in the limit of vanishing net baryon number density, it is necessary to extend the existing calculations of the equation of state and basic static properties of hot and dense matter in several respects: 1) extend scientists current knowledge of the equation of state to higher temperatures; 2) establish better theoretical control over the low temperature regime of the equation of state; and 3) better understand the dependence of thermodynamics on the light-quark masses to be able to explore the phase transition in the chiral limit of massless quarks.

Equation of State

Basic features of the temperature dependence of the energy density and pressure have already been established through lattice quantum chromodynamics (LQCD) calculations with rather crude approximations to continuum QCD. Calculations on “coarse” (large lattice-spacing) lattices with light-quark masses that are significantly larger than those of nature have shown that a change in the relevant degrees of freedom occurs over a narrow temperature interval (Karsch et al. 2001). However, even with the most current calculations (Bazavov et al. 2009), full control over the structure of the QCD equation of state has yet to be obtained. At high temperatures, contact has not been established with well-defined analytic calculations. At low temperature, the influence of chiral symmetry breaking and its impact upon the hadronic component of the equation of state has not been established. Moreover, the relevant degrees of freedom that control the structure of the equation of state in the transition region have not been determined. Is the restoration of chiral symmetry of any relevance to the QCD transition, or is the copious production of resonances the driving mechanism that leads to deconfinement and a strongly interacting medium of quarks and gluons at high temperature? To answer these questions, LQCD calculations of thermodynamic quantities at lower temperatures must be performed. In addition, lattice discretizations of QCD that respect chiral symmetry, or at least significantly reduce the influence of its explicit breaking due to the finite lattice spacing, are required in the transition region.

Chiral Fermions

To go beyond the current state-of-the art calculations of the QCD equation of state, it is necessary to use improved discretization schemes for the QCD action that respect all of the symmetries of the continuum theory. These discretization schemes have been developed over several years and continue to be

improved through further development. However, these discretization schemes have not been used extensively for numerical studies of QCD to date. This is because they require significantly larger computational resources to perform calculations with sufficiently small statistical uncertainties to allow for a meaningful comparison with the numerical results obtained with non-chiral discretizations.

While improved staggered fermion actions like the highly improved staggered quark (HISQ) (Follana et al. 2007) and stout (Morningstar and Peardon 2004) actions will be used extensively on petaflop computers, truly chiral formulations—such as domain wall and overlap fermion actions (Jansen 2008) — will require extreme scale computing resources in order for a comprehensive study of chiral aspects of the QCD equation of state. These discretized versions of the QCD action provide significantly better control over the chiral properties of QCD, and thus will be important for analyzing the low temperature and transition region of the static, bulk thermodynamic observables, for calculating hadronic screening lengths (Detmold and Savage 2009), and for determining order parameters that characterize the state of matter at high temperatures. Calculations with chiral fermions will enable the analysis of the universal properties of the transition, such as the scaling behavior of the chiral condensate, its susceptibility as well as quark number susceptibilities, and their fluctuations. Further, this work will provide a clarification of the relation between the QCD equation of state and the phenomena of deconfinement and chiral symmetry restoration.

High-Temperature Limit

Properties of strongly interacting matter at temperatures as large as three to four times the transition temperature will be probed experimentally at the Large Hadron Collider at the European Organization for Nuclear Research (CERN), Switzerland. At these temperatures, it may begin to be possible to make contact with perturbative calculations in finite temperature and density QCD (Kajantie et al. 2003; Vuorinen 2003). This will allow for a cross-check between numerical and analytic techniques used in this regime. A reliable numerical calculation of the equation of state and various screening lengths at such high temperatures requires large computational resources as large lattices are needed to control the renormalization of thermodynamic quantities through a proper subtraction of zero temperature observables. This allows for an elimination of otherwise divergent contributions that would prohibit a controlled extrapolation to the continuum limit. Recently developed techniques that minimize the required input from large zero-temperature calculations (Endrodi et al. 2007; Umeda et al. 2009) have the potential to make these calculations less demanding.

Computational Challenge

Calculations with domain wall fermions or overlap fermions require approximately two orders of magnitude more computational resources than calculations performed with staggered fermions. Prospects for the next generation of studies of bulk thermodynamics based on the staggered fermion discretization scheme have been examined in a white paper written in 2007 by the USQCD collaboration (USQCD 2007). This led to the conclusion that a thorough analysis of the equation of state at temperatures below twice the transition temperature will require approximately 100 sustained teraflop-years. Extending such a study to temperatures twice as high will increase the numerical effort by almost an order of magnitude. A thorough study of the QCD equation of state in the transition from low to high temperature needs to be performed with domain wall or overlap fermions. Such calculations require extreme scale computing resources, as shown in Figure 52.

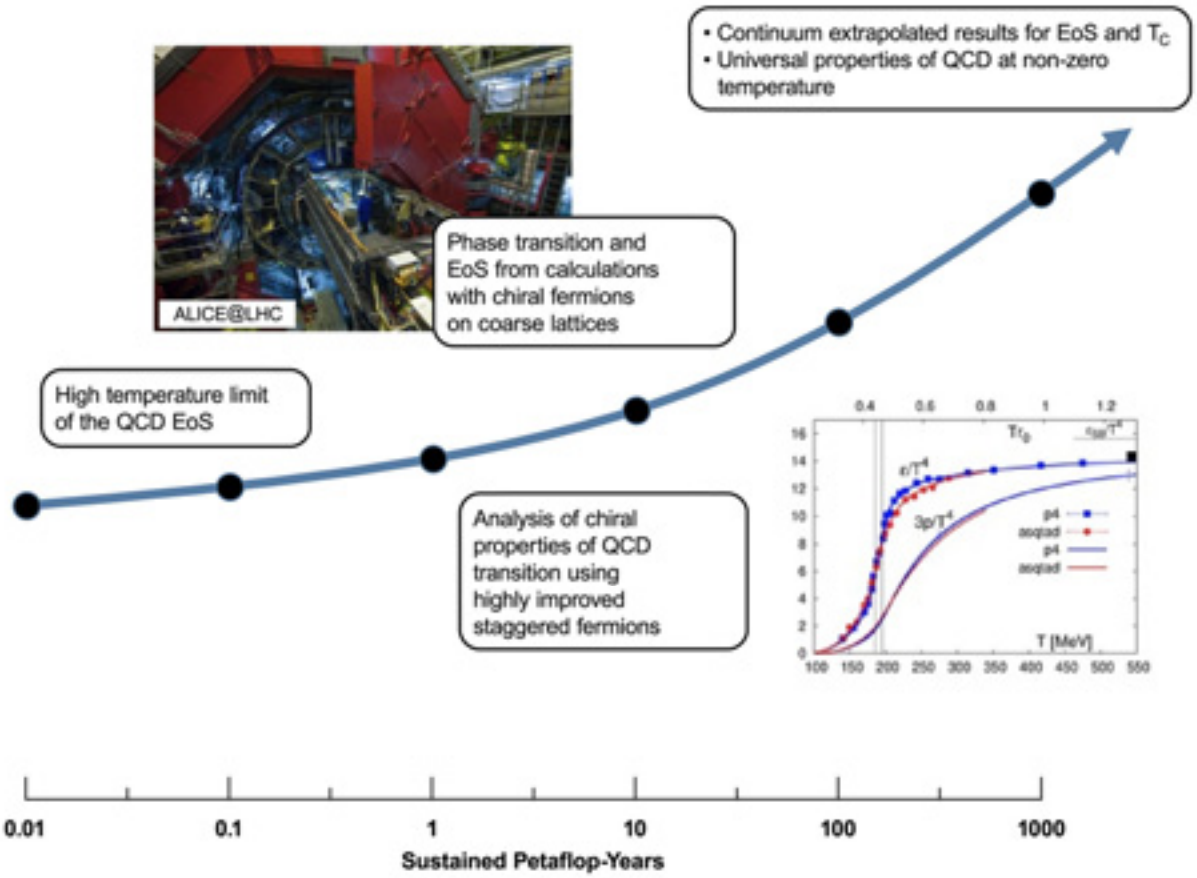

Figure 52. Anticipated highlights for priority research direction "Precision Calculation of Bulk Thermodynamics." Upper-left image courtesy of CERN. Remainder of image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

Scientific Outcomes and Impacts

Establishing the properties of strongly interacting matter at vanishing net baryon number density in the chiral limit will define the anchor point for all studies of the QCD phase diagram as a function of temperature and net baryon number density. It will establish a reliable starting point for extensions of these calculations into the regime of nonvanishing net baryon number density. In combination with calculations using values of light and heavy quark masses as realized in nature, this will quantify the role of chiral symmetry breaking and confinement in the thermodynamics of strongly interacting matter. The equation of state will be the basic equilibrium input to a microscopic description of the rapidly expanding and cooling dense matter formed in a heavy ion collision.

The calculation of the equation of state with physical values of the light-quark masses will not only have a significant impact on the modeling of heavy ion collisions, but it will also constrain the range of validity of conventional perturbative calculations at high temperatures and of model building, based on effective theories, at low temperatures.

QUANTUM CHROMODYNAMICS PHASE STRUCTURE AT NONZERO NET BARYON NUMBER DENSITY

Basic Scientific and Computational Challenges

Current studies of the QCD phase diagram and the thermodynamics at nonzero net baryon number density are limited to the region of small chemical potential; i.e., small net baryon number density. Sensitivity to possible phase transitions at larger values of the chemical potential could arise from conceptually new approaches to the LQCD calculations that overcome the sign problem. This might be achieved through the introduction of auxiliary degrees of freedom that eliminate the oscillating integrands in the QCD partition functions. The complex Langevin approach (Karsch and Wyld 1985; Aarts and Stamatescu 2008) may eventually lead to such an algorithm that avoids the sign problem. However, it has not yet been successfully implemented in realistic calculations. In the absence of such innovative concepts, currently explored techniques will need to be refined to perform calculations with substantially higher numerical accuracy. These numerical approaches include the Taylor expansion of thermodynamic quantities, such as the pressure and energy density, the analytic continuation of results from numerical calculations performed at imaginary baryon chemical potential, as well as approaches that allow for a projection onto physical states with a fixed baryon number. To use these methods in numerical calculations with physical parameters and improved discretization schemes is challenging and goes beyond currently performed exploratory studies.

Taylor Expansion Techniques

To extract sufficient information on the existence of phase transitions in the QCD phase diagram from a series expansion of the QCD partition function (Gavai and Gupta 2003; Allton et al. 2003*)*, which directly gives the expansion of the pressure as function of the baryon chemical potential, many expansion coefficients must be determined. This allows for a systematic analysis of the convergence properties of the series and provides insight into the analytic structure of the partition function. The required numerical effort grows rapidly with the order of expansion. Approximately two orders of magnitude increase in computing resources is required to calculate each additional nonvanishing order in the series expansion.

Analytic Continuation

A straightforward way to avoid the sign problem in calculations at nonvanishing net baryon number density is to replace the baryon chemical potential with a purely imaginary chemical potential (de Forcrand and Philipsen 2002; D'Elia and Lombardo 2003). This enables the use of the highly optimized algorithms developed for the calculation of the QCD equation of state at vanishing chemical potential. In particular, it is possible to perform calculations on large lattices with improved actions. However, to extract information on the thermodynamics at nonvanishing net baryon number density, extremely precise information is needed on the dependence of thermodynamic observables on the imaginary chemical potential. Only then is it possible to analytically continue (i.e., extrapolate) the numerical results to the physically relevant finite density regime.

Canonical Ensemble

An attractive, but extremely computationally demanding approach in the numerical studies of strongly interacting matter at nonzero net baryon number density, is to perform the calculations directly at a fixed

value of the net baryon number density (Kratochvila and de Forcrand 2005; Alexandru et al. 2005). This is in contrast to the approaches discussed above, where calculations are done with an auxiliary control parameter (the baryon chemical potential). To perform calculations in the so-called canonical ensemble generally requires the exact calculation of the determinants of large-sparse matrices. This is straightforward but computationally demanding. Such calculations may profit from improved eigenvalue solvers optimized for QCD applications.

Color Superconducting Phases

At low temperatures, but with large net baryon number density, QCD is predicted to become a color superconductor (Rajagopal and Wilczek 2000; Alford et al. 2008). There may exist several distinct phases, with competing patterns of light-quark flavor-color-spin-momentum pairings. The existence of such phases may have consequences for understanding the evolution of the early universe and the formation of compact stellar objects.

Very little is known from numerical calculations about the phase structure of strongly interacting matter in this regime (away from the extreme asymptotic limits). First-principles calculations in this regime are presently performed only in QCD-like models (Hands 2007). A direct study within QCD will require the development of new techniques that can manage or circumvent the sign problem. Extreme scale computing resources are required to explore such phases.

Computational Challenge

At present, calculations of Taylor expansions up to the third order in the squared baryon chemical potential require about 100 teraflop-years. Extending these expansions to the fifth order will require resources of 1 exaflop-year. To pursue calculations at these high orders, it is necessary to improve the numerical techniques used to calculate Taylor expansion coefficients. Improved techniques for the inversion of large, sparse matrices (deflation) and the optimization of random source vectors (dilution) are currently being tested and are expected to significantly expedite these calculations. The computational challenges that must be addressed in calculations with imaginary chemical potentials are similar. Quantitative studies of finite density QCD, and a decisive calculation that verifies or excludes the existence of a critical point in the QCD phase diagram, require extreme scale computing resources as shown in Figure 53.

Scientific Outcomes and Impacts

Calculations at nonvanishing net baryon number density will greatly advance current knowledge of the phase diagram of strongly interacting matter. High-precision calculations of high-order Taylor expansions, as well as accurate calculations with imaginary chemical potential, will provide information on the analytic structure of the QCD partition function. This may allow definitive statements about the density and temperature dependence of the thermodynamics of dense matter to be made, and eventually may determine the location (or may rule out its existence) of a critical point in the QCD phase diagram.

These calculations will have an enormous impact on current understanding of properties of strongly interacting matter. Further, they will provide strong constraints on the development of theoretical models for the high-density regime of strongly interacting matter, and will influence the accelerator-based experimental research program in this area.

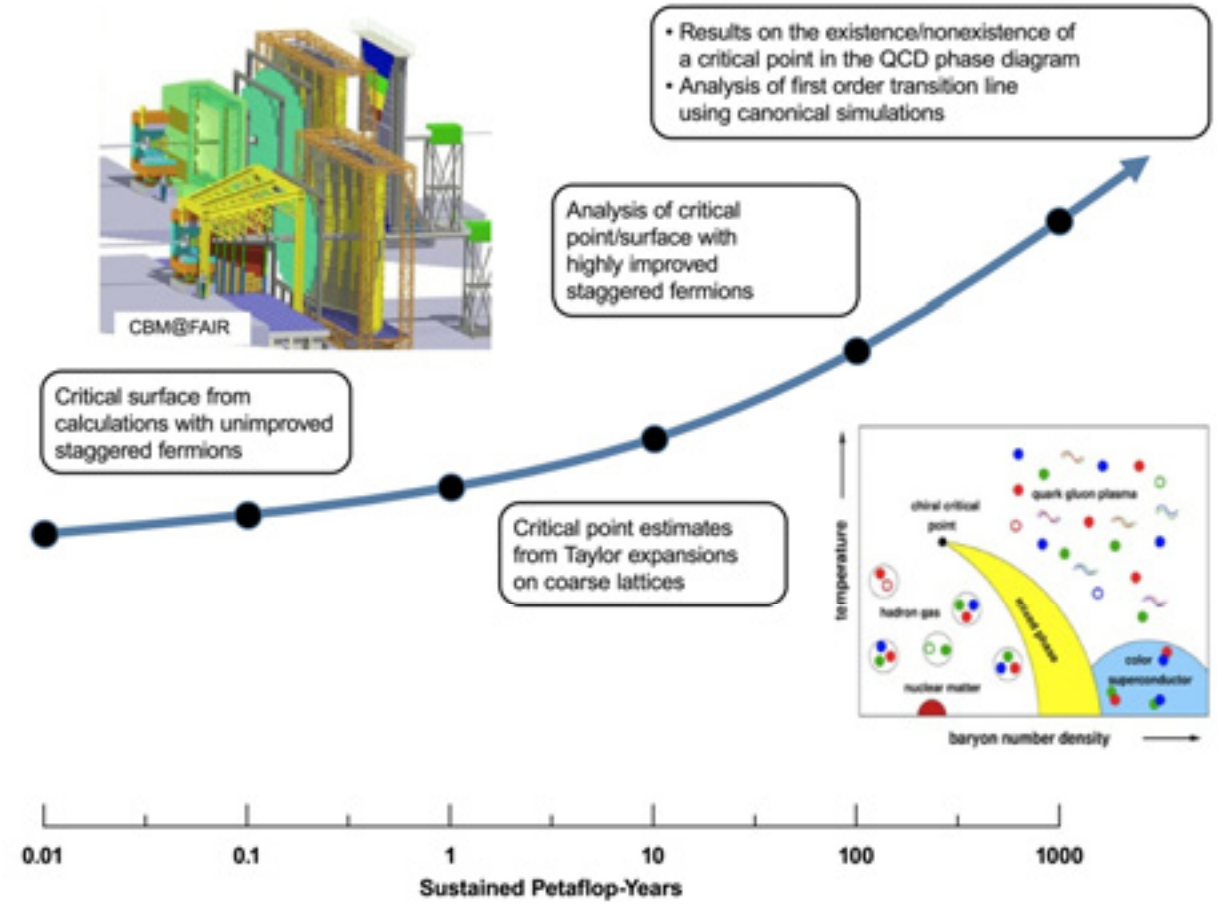

Figure 53. Anticipated highlights for priority research direction "Quantum Chromodynamics Phase Structure at Nonzero Baryon Density." Image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

TRANSPORT COEFFICIENTS OF QUANTUM CHROMODYNAMICS AND SPECTRAL FUNCTIONS OF HADRONS IN MEDIUM

Basic Science Challenges and Computational Challenges

Numerical calculations of the dynamic properties; i.e., the spectrum of excitations in hot and dense, strongly interacting matter, as well as transport properties of the medium, are presently performed at an exploratory level. To go beyond qualitative statements and to reach a point where quantitative predictions of dynamic properties become feasible, calculations on thermal lattices with unusually large spatial volumes, with greater than 10^3 times the number of lattice sites used in present day calculations, must be performed.

Transport Coefficients

The calculation of transport coefficients, such as the shear and bulk viscosity, that characterize the response of the medium-to-small deviations from its equilibrium state, are particularly difficult. Their calculations formally require taking the limit of zero frequency in an infinite spatial volume, which of

course, is not possible in numerical calculations. To obtain information on the excitation spectrum of a thermal medium requires accurate calculations of correlation functions at a large set of time separations. The extraction of the (continuous) spectral function from a finite set of data points is ill-posed (Karsch and Wyld 1987). To constrain the class of spectral functions that is consistent with these data, the noise level of the data set at the largest time separations has to be below the percent level. Furthermore, the correlation functions have to be calculated at a large number of time separations between sources, making use of correlations between different members of the data set. Thus far, such calculations have only been pursued in quenched QCD (Nakamura and Sakai 2005; Meyer 2007). Even in the quenched case, the lattices that were used were too small to obtain reliable results. A petaflop-year of computing resources will be required to complete the studies of transport properties in quenched QCD, and a computation with light dynamical quark degrees of freedom requires extreme scale computing.

In-Medium Hadron Masses

The degree of difficulty in calculating the hadronic excitations of the medium that provide information on the in-medium modification of light and heavy quark bound states, is similar to that of transport coefficients. Hadrons in a thermal bath interact with the medium, and these interactions can lead to the destruction of bound states, and thus the disappearance of the corresponding resonance peaks in the spectral function. Such an effect has been advocated as an experimental signature for the formation of a hot and dense medium in heavy ion collisions (Matsui and Satz 1986). Indeed, LQCD calculations of spectral functions at high temperature clearly demonstrate the disappearance of resonance peaks from the hadronic spectral functions (Nakahara et al. 1999). However, to follow the disappearance of these states in hot and dense matter in detail, and locate the melting temperature for various hadronic excitations, requires considerably more computing resources than are currently available. Prior to the disappearance of a state, interactions with the thermal medium will lead to temperature and density dependent shifts of the resonance peaks, as well as a broadening of these peaks. To resolve the structure of spectral functions to such a degree that shifts in resonance peaks and broadening of the spectral curve become statistically significant, accurate numerical results for hadron correlation functions are required. As in the case of calculations of transport coefficients, large lattices are needed to generate information on the correlation functions at many different time separations.

Computational Challenge

The major computational challenge in studies of the excitation spectrum of hot and dense matter is the quest for statistically accurate data on correlation functions on large lattices. These lattices are typically a factor of 50 larger than those used in calculations of static, bulk thermodynamics. The size of data samples needed to reach sufficiently small uncertainties in the correlation functions is approximately an order of magnitude larger. Fortunately, such calculations would only be performed at a few selected values of the temperature rather than at the large set of temperature values needed to control properties of the equation of state. Still, this presents a computational challenge, and requires a few petaflop-years to perform calculations within the quenched approximation to QCD. A fully dynamical LQCD calculation, which includes the light-quark contributions, will require extreme scale computing resources, as illustrated in Figure 54.

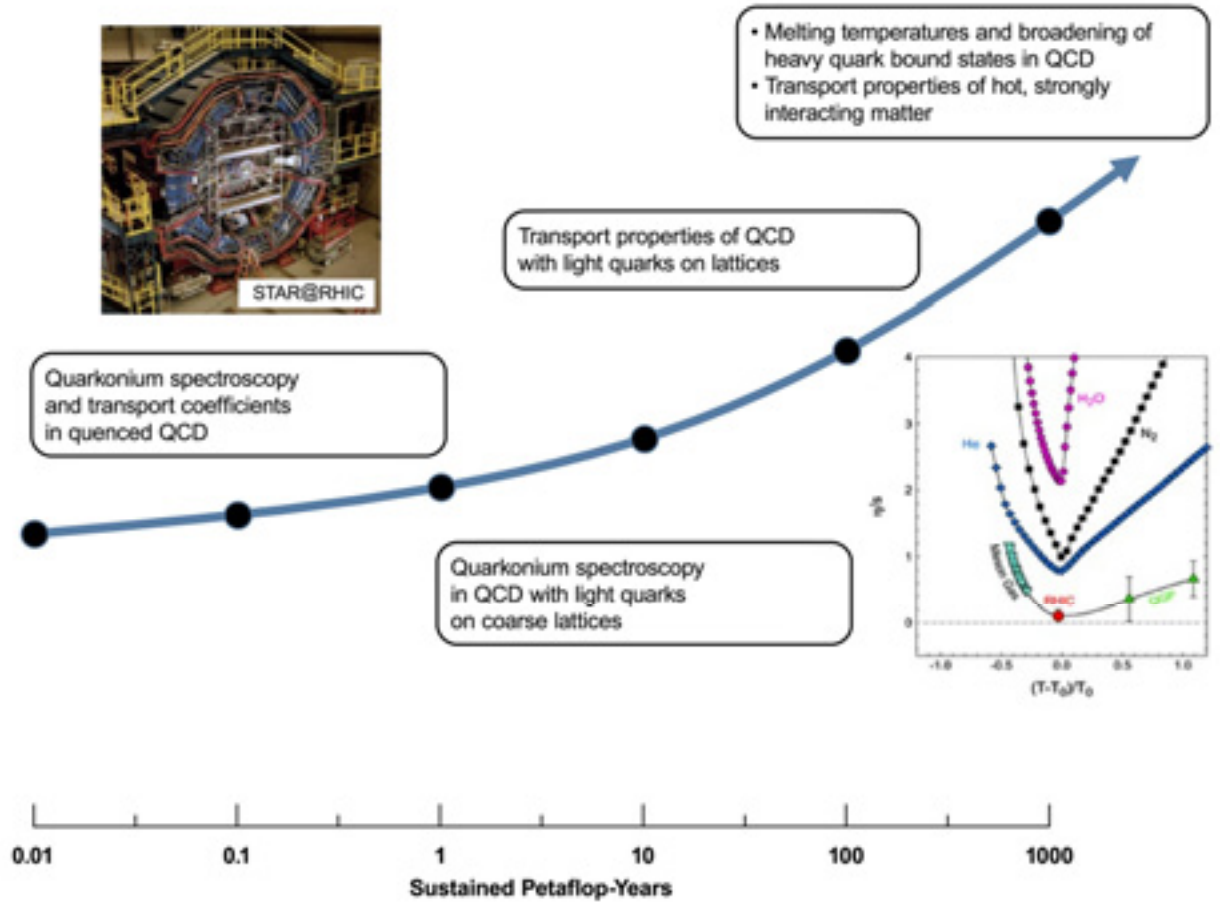

Figure 54. Anticipated highlights for priority research direction "Transport Coefficients of Quantum Chromodynamics and Spectral Functions of Hadrons in Medium." The point labeled RHIC in the right graphic is a theoretical "estimate." Image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

Scientific Outcomes and Impacts

Calculations of transport coefficients will provide fundamental insight into the structure of hot and dense matter. It will allow us to quantify aspects of the extent to which the phenomenologically successful modeling of heavy ion collisions has a solid foundation in QCD; i.e., whether a near-equilibrium quark-gluon plasma (QGP) described by QCD indeed equilibrates rapidly and can be characterized as an almost-perfect fluid. Detailed information on the spectral function would confirm whether or not the QGP is strongly coupled at the Relativistic Heavy Ion Collider (RHIC), and by varying the temperature in the LQCD calculations, scientists may learn how much the temperature has to be increased before the plasma becomes weakly coupled. This question will be of importance in comparing the heavy ion data obtained at the RHIC and the Large Hadron Collider experiments because the temperature in the latter will be about a factor 1.5 to 2 higher than in the former.

These calculations will strongly influence the analysis of experimental data obtained in heavy ion collisions.

EQUILIBRATION CHALLENGE: FROM THE COLOR GLASS CONDENSATE TO THE QUARK- GLUON PLASMA

Basic Science Challenges and Computational Challenges

Determining the time required for a QGP to form after the onset of the collision (i.e., the "thermalization" time) and determining the physics processes that drive the QGP formation are among the most important outstanding problems in the area of relativistic heavy ion collisions. The success of near-ideal hydrodynamics in describing bulk observables—such as the elliptic flow of matter created in noncentral collisions—implies that the matter has a short thermalization time compared to the overall timescale of the reaction.

To describe the approach to equilibrium, the following actions are required: 1) a firm understanding of the initial configuration of partons in the colliding nuclei and the process by which they are liberated from the nucleus at the onset of the collision needs to be acquired; and 2) detailed models and simulations of the processes that occur during the early nonequilibrium phase of the collision leading to the formation of a nearly thermalized QGP need to be developed.

Initial State of the Collision: Color-Glass-Condensate

A large nucleus moving near the speed of light contains a very dense system of gluons. It is believed that nonlinear effects in QCD lead to a saturation of the rapid growth of the gluon density in the colliding nuclei with beam energy and mass number A when the phase-space occupation number is (nonperturbatively) large, on the order of $1/\alpha_S$. The effective theory describing this nonlinear regime of QCD is the color-glass condensate (CGC) (McLerran and Venugopalan 1994a, 1994b; Kovchegov 1996). McLerran and Venugopalan (1994a) proposed an effective action incorporating high-gluon density effects, which amounts to solving the classical Yang-Mills equations where the large-momentum degrees of freedom in the nucleus act as sources of color charge for the small-momentum degrees of freedom.

The classical description of gluon saturation is modified at higher energies due to quantum loop corrections. To this end, a new set of equations, commonly referred to as the JIMWLK (Jalilian-Marian – Iancu – McLerran – Weigert – Leonidov – Kovner) equations (Jalilian-Marian 1997a, 1997b; Iancu et al. 2001), have been derived from a Wilsonian Renormalization Group formalism. They describe the evolution of n-point functions in QCD with energy. The resulting equations are an infinite hierarchy of coupled differential equations that are difficult to solve analytically. Nevertheless, they can be written in a form which, in principle, allows them to be solved by lattice gauge-theory techniques (Rummukainen and Weigert 2004).

The CGC has explained several theoretical and phenomenological aspects of high-energy interactions quite successfully. Nevertheless, many important properties of the CGC remain to be addressed quantitatively and tested by comparison with experimental data from RHIC and other colliders.

Thermalization Mechanisms: Plasma Turbulence

In recent years, it has been shown that early studies of the driving mechanism for the equilibration of quark-gluon matter in ultra-relativistic heavy ion collisions overlooked a crucial aspect of the dynamics of nonequilibrium plasmas—namely, the possibility of plasma instabilities. Most importantly, it has been shown that these instabilities may produce plasma isotropization and approximate thermalization on time-scales relevant to relativistic heavy ion collisions. The possibility of such non-Abelian plasma

instabilities was first predicted in the mid-1990s (Mrówczynski 1993) by studying plasmas with an anisotropic momentum-space distribution. The resulting instability has been dubbed the chromo-Weibel instability. In recent years, this theory has received a significant amount of attention because of analytic and numerical advances. The first advance was to show the instabilities predicted by Mrówczynski are generic and independent of the precise details of the assumed anisotropic parton distribution function (Arnold et al. 2003, 2004; Romatschke and Strickland 2003). Various schemes for treating the non-Abelian, nonequilibrium dynamics (Mrówczynski et al. 2004) using real-time lattice gauge techniques (Hu and Müller 1997) are being explored.

Basic Scientific and Computational Challenges

The full description of the initial state of a heavy ion collision and subsequent thermalization of the matter requires code that can self-consistently describe both the earliest periods when the physics of saturation is important, and also the intermediate times when the physics of the chromo-Weibel instability becomes important. To do this requires the real-time solution of the Yang-Mills equation on three-dimensional lattices coupled self-consistently to the Wong equations. Such codes already exist for simplified configurations and expansion scenarios (Bass et al. 1999; Dumitru et al. 2007; Schenke et al. 2008). The solution of the full problem requires three-dimensional lattices with a fine lattice-spacing in the longitudinal direction; i.e., solving the classical Yang-Mills equations in real time on a three-dimensional lattice with approximately 512^3 sites. The field equations describing the low-momentum gluons need to be coupled self-consistently to the Wong equations that describe the propagation of the hard valence sources in the soft background, including energy-momentum conservation. Beyond the classical limit, a simultaneous solution of the rapidity dependence of the JIMWLK measure together with the real-time evolution of the initial fields is required. Beyond that, it is necessary to have a lattice that is capable of describing the dynamics of the chromo-Weibel instability and the subsequent non-Abelian cascade to high-momentum modes. The "brute force" method to accomplish this is to ensure the lattice spacing is sufficiently fine. This means using lattices significantly larger than 512^3. As some of the processes (e.g., initial conditions, binary particle collisions, and hard radiation) are stochastic, it will be necessary to average observables over multiple sets of initial conditions (runs). Currently, each run of the simplified calculation requires approximately one teraflop-year. Factoring in the higher dimensionality required for the full problem increases this estimate into the tens of petaflop-years region. Averaging over initial conditions and varying experimental parameters will only be possible with extreme scale computing resources as illustrated in Figure 55.

Scientific Outcomes and Impacts

Extreme scale computing will deliver real-time calculations of the collision of two heavy ions at high energy, retaining the complete three-dimensional structure of the fields of produced gluons (in impact parameter and rapidity space), energy-momentum conservation, and quantum evolution of the measure. The subsequent real-time evolution of the color fields following the initial impact will clarify the timescales and processes that lead to thermalization and formation of a QGP and the possible role played by non-Abelian gauge-field instabilities (plasma-turbulence). The distribution of the thermalized gluons in the impact parameter and rapidity will provide much-needed initial conditions for hydrodynamic modeling of the late stages of the collision, and could provide information on the equation of state and the viscosity of hot QCD matter. This work will also provide predictions for the effect of early-time nonequilibrium dynamics on important QGP observables such as jet quenching, anomalous transport, and fluctuations.

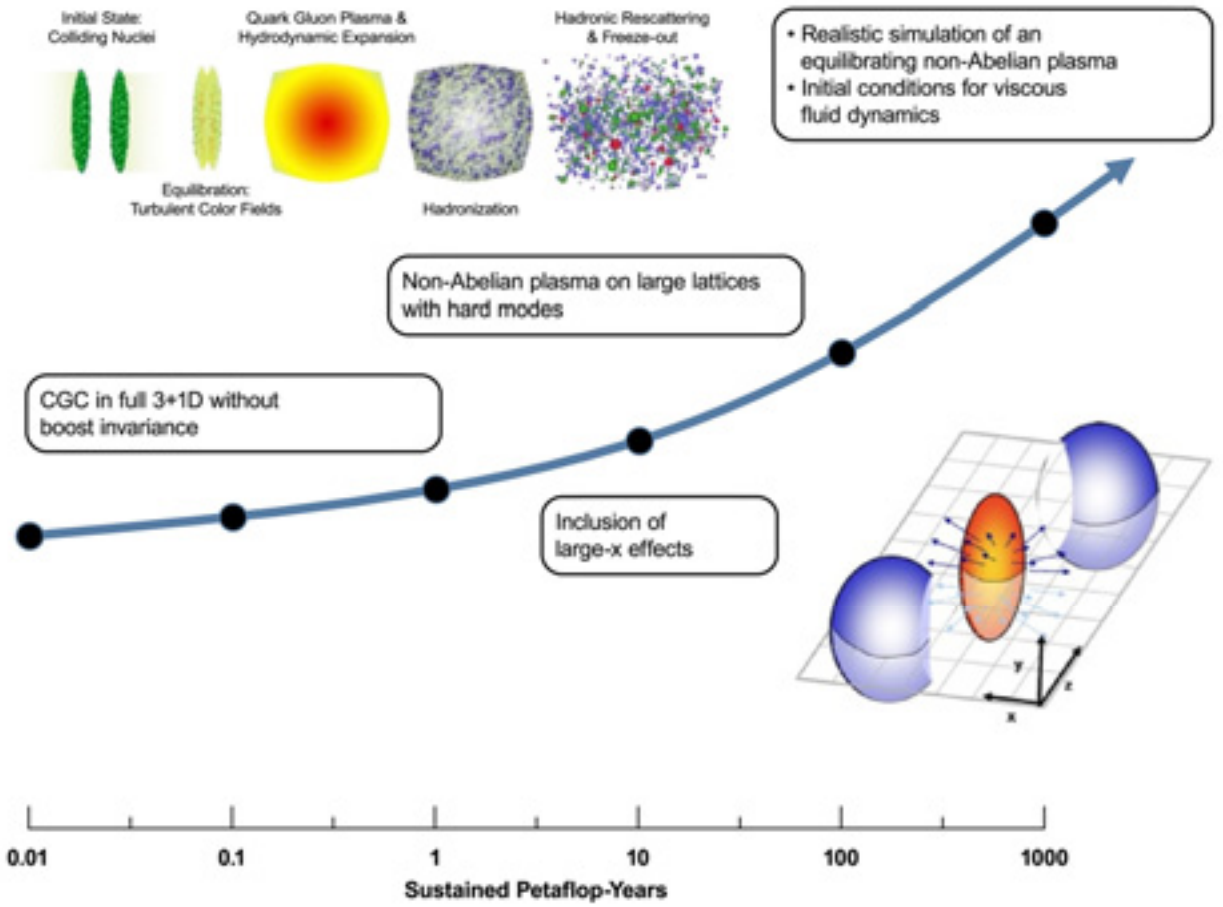

Figure 55. Anticipated highlights for the priority research direction "Equilibration Challenge: From the Color Glass Condensate to the Quark-Gluon Plasma." Lower-right image courtesy of Jerome Lauret. Remainder of image courtesy of Steffen A. Bass (Duke University) and Frithjof Karsch (Brookhaven National Laboratory).

ACCELERATOR PHYSICS

MAXIMIZE PRODUCTION EFFICIENCY, VARIETY, AND PURITY OF RARE ISOTOPE BEAMS FOR NUCLEAR PHYSICS EXPERIMENTS

Basic Scientific and Computational Challenges

The design of the Facility for Rare Isotope Beams (FRIB) poses several challenges. These include the development of new techniques for optimal design and tuning of isotope separators to select and to purify extremely rare isotopes, advances in electron cyclotron resonance (ECR) on source modeling, advances in radio frequency quadrupole (RFQ) modeling, optimal design of low-beta radio frequency cavities, and advances in beam-dynamics modeling and optimization offline and in near real time. Figure 56 highlights several of the computational challenges associated with the design of FRIB, starting at the petascale and ranging to the extreme scale.

Optimal Design and Tuning of Isotope Separators

One of the most important aspects of experimenting with rare isotopes is related to the separation purity of the isotopes to be selected. If the reaction mechanism involves fragmentation and/or fission of heavy ion primary beams, then a multistage isotope separator is employed. The separator must be designed such that the primary beam rejection is perfect, and the several hundreds of other species that act as unwanted impurities for the experimental stations are minimized. The tiny production cross-sections of rare isotopes of interest to nuclear physics make this an especially difficult problem to model and simulate accurately and efficiently. The computational challenges are dependent on the rarity of the isotopes to be studied. These challenges include research currently possible on a typical single-processor personal computer (such as ^{14}Be obtained from fragmentation of oxygen requiring approximately 10-gigaflop-days), cases that are manageable with petaflop machines (such as ^{132}Sn shown in Figure 56 requiring approximately 1 petaflop-day), and cases that require extreme scale computing resources (such as ^{100}Sn from fragmentation of xenon requiring approximately 100 petaflop-years).

All cases noted above become extreme scale computing problems when design or experimental setup optimization is included. Global parallel-parameter optimization (magnet strengths, slit settings, target and wedge materials and thicknesses, and absorber shapes) is a challenge in its own, so the computational resource requirement is significantly increased with respect to that stated in the previous paragraph if these experimental parameters are added to the problem (factors of thousands more than above). These applications therefore need extreme scale computing resources. Some exceptionally difficult cases may require of the order of 100-exaflop-years.

The computational challenge may be even greater depending on the physics of the problem and the complexity of the parallel optimization model. Coupling to radiation transport codes, inclusion of space charge effects, and inclusion of full electromagnetic design in the optimization loop extend the problem to the extreme scale and beyond.

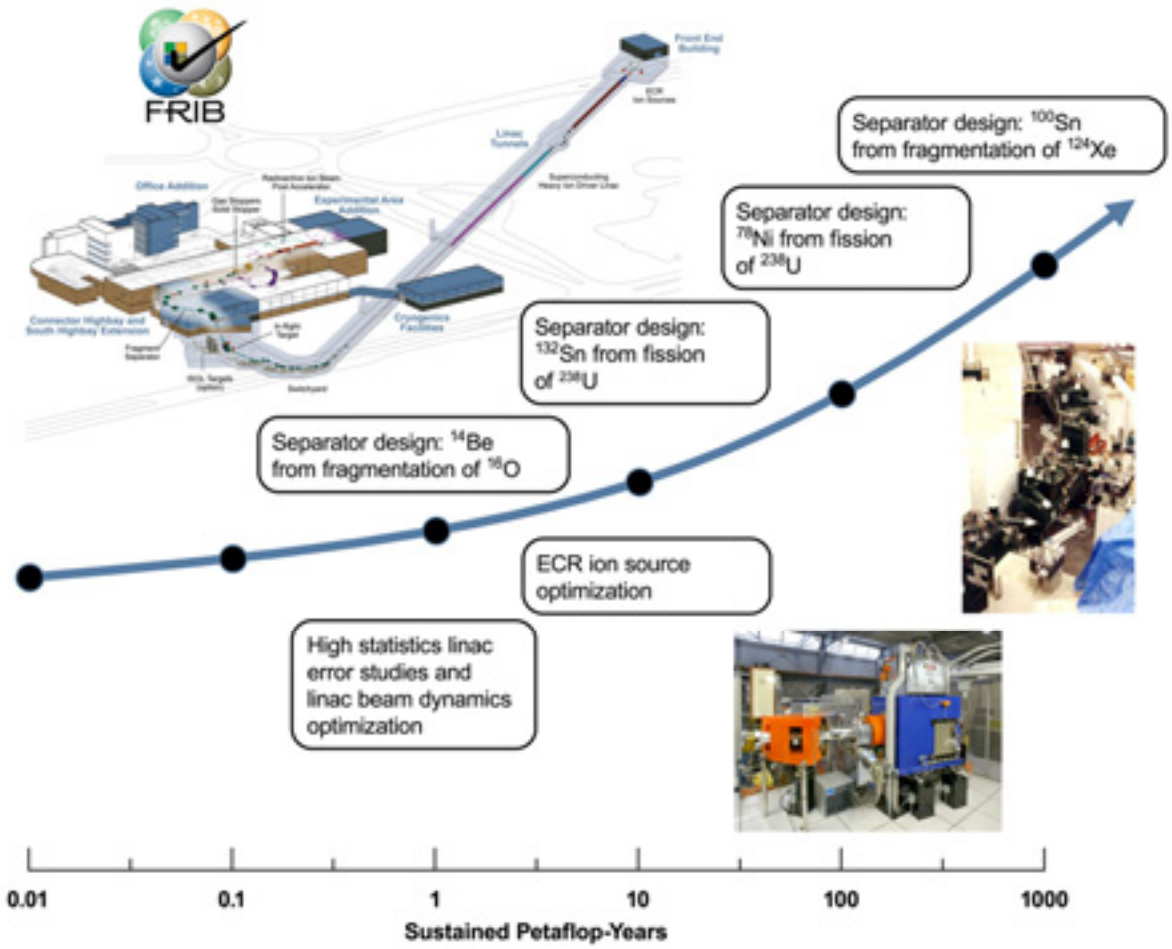

Figure 56. Anticipated highlights for priority research direction "Maximize Production Efficiency, Variety, and Purity of Rare Isotope Beams for Nuclear Physics Experiments." Upper-left image and right images courtesy of Michigan State University. Bottom-right image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

Advances in Electron Cyclotron Resonance Ion Source Modeling

For rare isotope facilities—such as the recently sited FRIB—ECR ion sources are required to produce a wide variety of highly ionized isotopes. These sources provide high currents of multiply charged ions for injection in the main driver accelerator. Large currents are needed for the physics program, and high charge states are needed to make most efficient use of available accelerating voltage.

Modeling ECR ion sources, including the ion beam extraction, is critical for better understanding and subsequent optimization of these sources. To date, reduced models have not agreed with experimental data. It is therefore necessary to use electrostatic and electromagnetic particle-in-cell (PIC) simulations that include multiple physical phenomena such as impact ionization, charge exchange, recombination, and particle-wall interactions. However, the extreme variation of spatial and temporal scales has limited such efforts to artificially small domains and short times (Mullowney et al. 2008). Emerging leadership-class computational facilities will enable full-scale, first-principle simulations.

The VENUS ECR source (Venus ECR) is ~1 m across, while the electron Debye length, λ_D, is ~10 microns, and the complicated magnetic fields (solenoid plus sextupole) are fully three-dimensional. The cell size can be chosen larger than λ_D if high-order particle shapes are used to prevent numerical heating. For $\Delta x \sim 50\lambda_D$, a mesh with $\sim 10^{10}$ cells is required, and a large fraction of the domain will have ~ 50 particles per cell (electrons plus multiple ion species), leading to $\sim 3 \times 10^{11}$ macroparticles. The timescale for ion dynamics (i.e., confinement, extraction, etc.) is several milliseconds, while the time step must be small enough to resolve the ionization/recombination timescale of 1 microsecond, which implies $\sim 10^5$ time steps.

For high-order particle shapes with multiple collisional processes, approximately 10 core-micro seconds are required for each particle step. Given a required number of 3×10^{16} particle steps in the calculation, and assuming efficient scaling up to approximately 1 million cores, this implies approximately 10^8 core-hour simulations, which translates into running at approximately 100 petaflop-hours. Given the need for a large number of such simulations (~1,000 per year) to design future and optimize existing ECR ion sources, the need for approximately 10 petaflop-years of computing resources can be predicted for each calendar year.

Advances in Electromagnetic Modeling and Optimization

Along with the above-mentioned advances in beam-dynamics modeling, advances in electromagnetic modeling and optimization that will be possible on extreme scale computers will have a major impact on the design of the FRIB. Examples include the design of ECR ion sources, of the FRIB RFQ, and of low-beta structures. These topics are described in more detail in the section titled, "Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities."

Advances in Beam-Dynamics Modeling and Optimization

Currently, several highly parallelized beam-dynamics codes (e.g., TRACK and IMPACT) are available for the simulation of the FRIB accelerators. The majority of the accelerator design work can be performed using petascale computers, as demonstrated by teams from Argonne National Laboratory, Michigan State University, and Lawrence Berkeley National Laboratory. The availability of extreme scale resources to support the commissioning and operation of the FRIB accelerators is essential to expedite the delivery, shorten the commissioning phase, and effectively operate the machine by applying a "model-driven accelerator" concept. Meanwhile, current parallel beam-dynamics codes need to be updated with large-scale optimization tools and must include beam diagnostics and control systems into the optimization loop. These are described in the section titled, "Advanced Methods and Applications of Accelerator Simulation for Nuclear Physics Facilities."

Scientific Outcomes and Impacts

The availability of extreme scale computing resources, along with accelerator codes that make use of this computer power, will significantly impact the design and operation of FRIB, as well as other U.S. Department of Energy Office of Science (DOE SC) accelerator facilities. In regard to FRIB, the availability of these resources will provide a means to design and operate isotope-separation systems with significantly increased accuracy and efficiency—and in some cases, will make seemingly impossible design simulations feasible. More generally, extreme scale computational resources will lead to optimized designs better able to meet facility requirements while reducing cost and risk. Beyond the

design phase, extreme scale resources, applied in the model-driven accelerator paradigm, will permit faster commissioning, improved diagnosis of operational issues, and improved facility operations. The successful design, commissioning, and operation of the FRIB, facilitated by large-scale and extreme scale computing, will lead to numerous advances in nuclear science. These include providing a comprehensive description of nuclei, elucidation of the origin of the elements in the cosmos, providing an understanding of matter in the crust of a neutron star, and establishment of a scientific foundation for innovate applications of nuclear science (DOE 2007).

DEVELOP OPTIMAL DESIGN FOR AN ELECTRON-ION COLLIDER

Basic Science Challenges and Computational Challenges

The design of a future electron-ion collider (EIC) poses several challenges. To address these challenges, scientists focus on the development of advanced computational capabilities for multiphysics accelerator modeling, and the development of advanced electron-cooling systems and of energy-recovery linac technologies. Figure 57 highlights several of the computational challenges associated with the design of an EIC, starting at the terascale and ranging into the extreme scale.

Design of Electron Cooling Systems

Cooling of the hadron beam in a future EIC is critical to meet performance requirements needed for scientific discovery at such a facility. At present, electron cooling is the most promising concept. An electron cooling system has not yet been operated in the regime of a proposed EIC. Hence, accurate physical modeling is essential. Previous work based on molecular dynamics in a small, idealized domain needs to be extended to a molecular dynamics/PIC approach that covers the full extent of the overlapping electron and ion beams over many Debye lengths. This results in a multiscale simulation that requires extreme scale computational resources. For beam and system parameters relevant to the Electron Relativistic Heavy Ion Collider (eRHIC) concept, a single simulation using VORPAL operating with approximately 10% of peak performance requires approximately 10 teraflop-hours. To design an electron-cooling system, scientists must run approximately 10^4 seeds simultaneously to obtain enough dynamical friction force values to adequately characterize the performance for a single set of electron parameters. This implies a resource requirement of 100 petaflop-hours. These runs would need to be repeated for a variety of electron-beam parameters and also magnetic fields (solenoid and/or undulator) to optimize the design. For an optimal case, many instances would need to be run with different sets of magnetic field errors to evaluate the corresponding reduction in the cooling strength. Over the course of a year-long design effort, one could easily expect to complete approximately 1000 such cases, leading to resource usage of approximately 10 petaflop-years over the course of one calendar year.

Energy-Recovery Linac Technologies

The energy-recovery linac (ERL) is an important emerging concept for building energy efficient and therefore less expensive accelerators. Accelerator physics researchers are considering including ERLs in electron-cooling sections and EICs for nuclear physics projects. One important issue for understanding ERLs is beam loss due to beam halo. On the recovery pass, the beam phase space is significantly modified (e.g., due to energy extracted as part of the lasing or cooling processes); thus, nonlinear and nonideal behavior may result in increased beam loss. Modeling beams with sufficient computational particles to resolve beam halo and beam loss is a significant challenge that requires extreme scale computing.

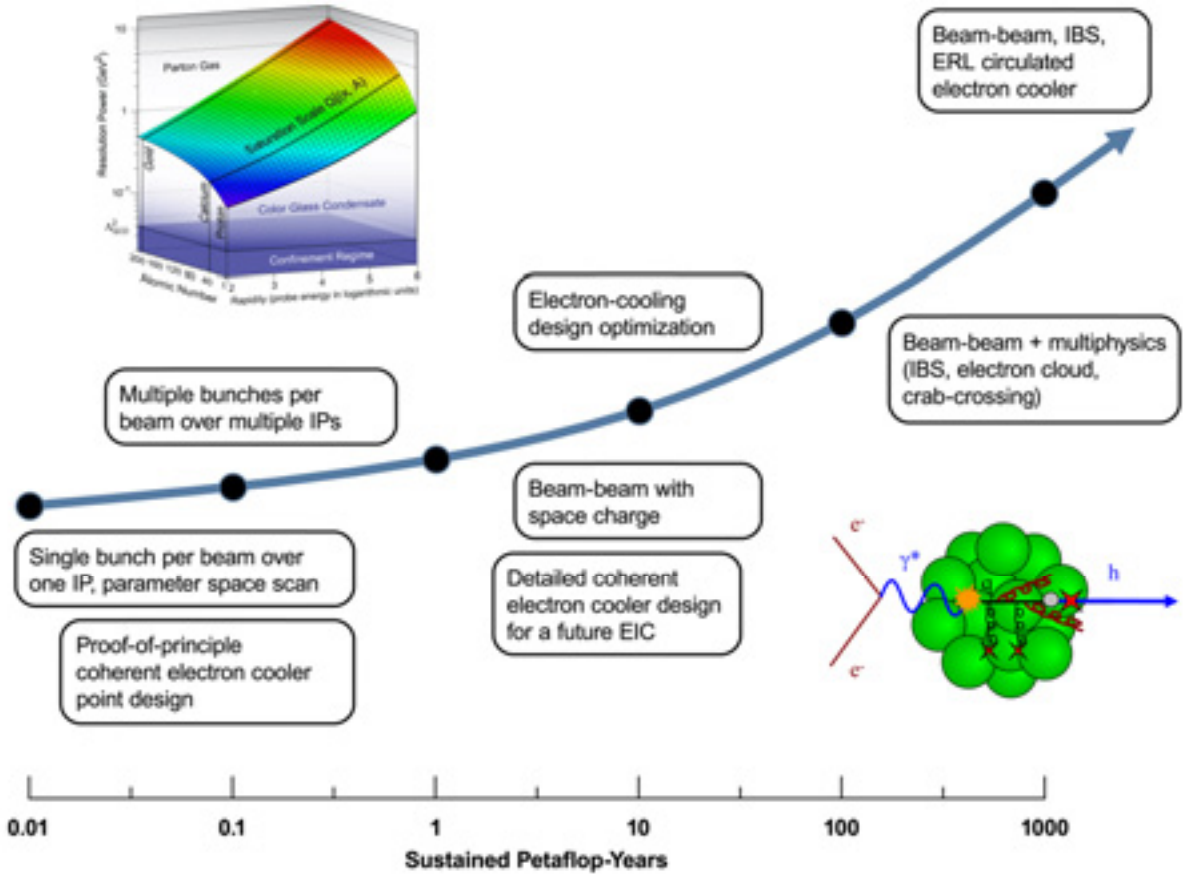

Figure 57. Anticipated highlights for priority research direction "Develop Optimal Design for an Electron-Ion Collider." Top-left image courtesy of Thomas Ullrich (Brookhaven National Laboratory). Bottom-right image courtesy of Alberto Accardi (Hampton University and Thomas Jefferson National Laboratory). Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

A second important issue related to ERLs includes multipass beam breakup and coupling with higher-order modes. Again, because the beam on the recovery pass has a modified and nonideal phase space, the beam may couple more strongly with higher-order modes—especially at the higher currents considered for future machines. These modes may feed back and produce beam breakup. Modeling beam breakup and coupling with higher-order modes will involve grid resolutions and simulation times that require extreme scale computing resources.

Finally, another important issue in ERL development is electron sources, where higher currents will produce issues of dark current, field emission, and cathode degradation due to ion bombardment. These higher currents will also produce subsequent heating issues associated with these effects. Modeling these effects in electron sources will require multiphysics code coupling and grid resolution that require extreme scale computing resources.

Scientific Outcomes and Impacts

The EIC will explore a new cold quantum chromodynamics frontier of strong color fields in nuclei, precisely image the gluons in the proton, and look for gluon saturation in heavier nuclei (DOE 2007). Large-scale, parallel computing is essential to design an EIC that maximizes luminosity and the potential for physics discovery. Beyond the development of an optimized EIC design, the design effort itself will lead to important new concepts and technologies for particle accelerators including new beam-cooling techniques, ERL technologies, advanced high-brightness sources, and new beam-beam compensation schemes.

DESIGN OPTIMIZATION OF COMPLEX ELECTROMAGNETIC STRUCTURES FOR NUCLEAR PHYSICS ACCELERATOR FACILITIES

Basic Science Challenges and Computational Challenges

Large-scale, high-accuracy electromagnetic design and optimization is essential to all major DOE SC accelerator projects. Several of the challenges and research directions have already been described in the "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale" workshop.[1] Examples include the development of novel, scalable eigen- and linear solvers for eigensystems at the extreme scale, as well as the implementation of efficient parallel algorithms for extreme scale modeling in a multiphysics environment. While many of the challenges and research directions are common to the various program offices of the DOE SC, each program office also has some unique needs and priorities. Evaluation of high-order mode (HOM) heating in the cryomodule of an ERL; design of the FRIB RFQ; and the multiphysics design and optimization of low beta radiofrequency accelerating cavities for FRIB—three priorities of the Nuclear Physics program—are described below. Figure 58 highlights the computational challenges associated with these areas, ranging from the terascale to the extreme scale.

High-Order Mode Heating in Cryomodule of Energy-Recovery Linac

One potential accelerator issue with ERL is HOM heating in the superconducting radio frequency cavities, especially with the combination of high current and short bunch length. It is important to estimate the broadband HOM power generated by the transit of the bunch to adequately dampen HOM without affecting cryogenic efficiency. To a great extent, the computational requirements are determined by the bunch length, whose frequency spectrum must be resolved by a fine mesh. Thus, bunch length and mesh density are inversely related. In typical nuclear physics applications, such as the ERL used in the proposed luminosity upgrade of RHIC known as RHIC-II, the bunch length is relatively long (roughly 1 cm). The electromagnetic simulation of broadband HOM in a cryomodule involves solving a linear system with 150 million degrees of freedom for 100,000 time steps. The estimated computational requirement is approximately 1 petaflop-hour, which can be handled easily by a petascale computer.

[1] "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale," December 9-11, 2008, Menlo Park, California. Additional information available at http://extremecomputing.labworks.org/highenergyphysics/index.stm.

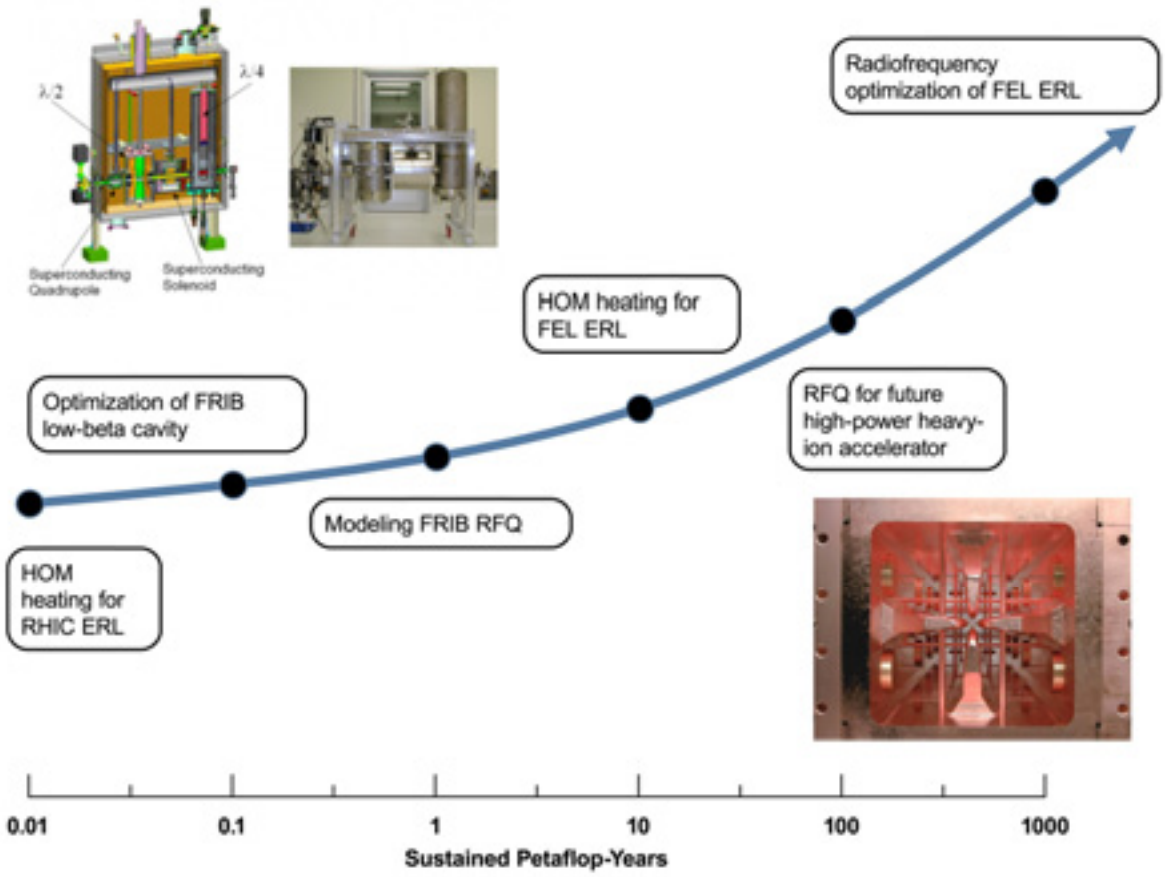

Figure 58. Anticipated highlights for priority research direction "Design Optimization of Complex Electromagnetic Structures for Nuclear Physics Accelerator Facilities." Upper-left images courtesy of Michigan State University. Lower-right image courtesy of Lawrence Berkeley National Laboratory. Remainder of image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

In typical Free-Electron Laser applications, such as the ERL in the Thomas Jefferson National Accelerator Facility infrared Free-Electron Laser upgrade, the bunch length is much shorter (roughly 1 mm or less) than at RHIC-II. The electromagnetic simulation of broadband HOM in a cryomodule involves solving a linear system with 150 billion degrees of freedom for 10^6 time steps. The estimated computational requirement is approximately 1 petaflop-year.

For a cavity system on the order of 10 cm by 100 cm, a simulation with a resolution on the order of microns would require 10^{14} cells. Simulating many tens of oscillations will require 10^7 time steps, implying that 10^{21} cell steps are required. For electromagnetic PICs with a finite difference time domain algorithm on current standard architecture, such as the Cray XT4^{TM}[1], this implies 10^{14} - 10^{15} core-seconds. Therefore, to complete a simulation in a single day would require 10^9 - 10^{10} cores, or a computing resource of approximately 3 petaflop-years.

[1] Cray XT4 is a trademark of Cray Inc.

Advances in Radio-Frequency Quadruple Modeling

The initial acceleration in FRIB is provided by an RFQ capable of accelerating any ion from hydrogen to uranium. The transport, bunching and focusing of charged particles in the RFQ resonator, is provided by appropriate design geometry of four modulated vanes. Parallel electromagnetic design tools with a fine-grained mesh must be used to develop a detailed resonator design for the RFQ. Accurate field maps of accelerating modes of cavities in the FRIB linear accelerator are required to reliably track particles in beam-dynamics studies. The long RFQ used in the low-energy linac could have disparate spatial scales ranging from fine transverse variations of the order of millimeters to longitudinal modulations of the order of a meter. While the frequency of the accelerator mode converges quickly with the reduction in finite element size, the field accuracy must be obtained with a much denser mesh or with the use of higher-order finite elements. The electromagnetic eigenmode simulation of the accelerating mode in the RFQ requires the solution of a linear system with 100 billion degrees of freedom. The estimated computational requirement is approximately 1 petaflop-day. This work is needed for future accelerators and accelerator upgrades, which will place greater demands on an injector than does FRIB, which will use the current state-of-the-art injectors.

Advances in Low-Beta Radio Frequency Cavity Design

The accelerator systems for the proposed FRIB will be based on several types of normal-conducting and superconducting transverse electric and magnetic mode (TEM-class) radio frequency cavities. The TEM-class superconducting cavities are much more efficient in the low-velocity (beta < 0.7) region compared to the elliptical superconducting cavities widely used for relativistic particles. Therefore, four to five types of TEM-class cavities are required in the FRIB driver linac. The integrated simulation of the resonators electromagnetic, thermal, and mechanical properties is necessary due to the complexity of the cavity's geometry. The optimal electrodynamic, thermal, and mechanical design of the fully dressed superconducting cavity together with fast tuner, slow tuner, radio-frequency coupler, and helium vessel could save millions of dollars. Obviously, the optimal design of the cavity should be free from multipacting. Currently, multiphysics design of a cavity can be performed using separate commercial software and does not allow researchers to provide full design optimization. Particularly, no code is available to study microphonics, which is due to coupled electro-mechanical oscillations induced by mechanical background noise. These simulations require a very large number of mesh points calling for extreme scale computing and a highly scalable multiphysics code.

Scientific Outcomes and Impacts

The development of an extreme scale, multiphysics cavity design and optimization tool will benefit not only the Nuclear Physics program but also the entire DOE accelerator complex. The process of developing these tools will also lead to advances in computational-enabling technologies such as linear solvers, eigensolvers, meshing, adaptive refinement, data analysis, and visualization.

Extreme scale electromagnetic modeling will shorten the design and build cycle for accelerator structures and components. It will also lead to cost savings by allowing researchers to explore and optimize designs that reduce cost while satisfying beam quality and machine operational reliability requirements. Lastly, the new modeling tools will provide accurate field maps of accelerating cavity modes for beam-dynamics simulation in overall machine studies.

ADVANCED METHODS AND APPLICATIONS OF ACCELERATOR SIMULATIONS FOR NUCLEAR PHYSICS FACILITIES

Basic Science Challenges and Computational Challenges

Extreme scale computing will dramatically change the way researchers design accelerators. It will enable the exploration of new accelerator concepts and the study of important phenomena, for which modeling was previously thought to be too computationally challenging. Extreme scale computing will also dramatically affect how scientists optimize accelerator designs. Lastly, extreme scale computing will change the way scientists' commission and operate accelerators.

Exploration of Advanced Concepts

Extreme scale computing will significantly impact the ability to explore innovative accelerators. The Accelerator Physics working group in the "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale" workshop identified the design of an ultracompact plasma-based collider as one of its key research directions.[1] The exploration and optimization of laser- and plasma-based concepts using extreme scale resources will have applications across the DOE SC. The design of fixed-field alternating gradient accelerators and the design of coherent electron cooling systems is discussed in the following sections.

The broad class of fixed-field alternating gradient accelerators is experiencing an international revival in the quest for high-beam power, duty cycle, reliability, and, in the case of the spiral-sector fixed-field alternating gradient, the potential for compactness at reasonable cost (Prior 2007; Johnstone and Koscielniak 2008, and references therein). The proposed fixed-field alternating gradients have the high average current and duty cycle characteristic of the cyclotron combined with the smaller aperture, beam losses, and energy variability of the synchrotron.

Because of frequently challenging field models, the computation of guiding, focusing, and accelerating fields and the assessment of stability of orbits usually must happen in an integrated mode of computation. Specifically, closed orbits and transfer maps around the orbits must be computed to a high order and in fine steps of up to 100 locations in the device. For each of these locations, it is common to perform a resonance analysis through normal form tools (Berz 1999) and orbit tracking of usually around 10^5 revolutions to assess stability. All these must be subjected to extensive design optimizations to arrive at viable machines. An example of one such simulation that illustrates the need for high-order models is shown in Figure 59. Currently, a sufficiently detailed field-design simulation and subsequent orbit analysis typically takes on the order of hundreds or thousands of core-hours—and it is expected that a future exhaustive automated search of parameter space may require on the order of 10^6 such iterations, leading to an overall cost in the range of exaflop-months.

[1] "Scientific Challenges for Understanding the Quantum Universe and the Role of Computing at Extreme Scale," December 9-11, 2008, Menlo Park, California. Additional information available at http://extremecomputing.labworks.org/highenergyphysics/index.stm.

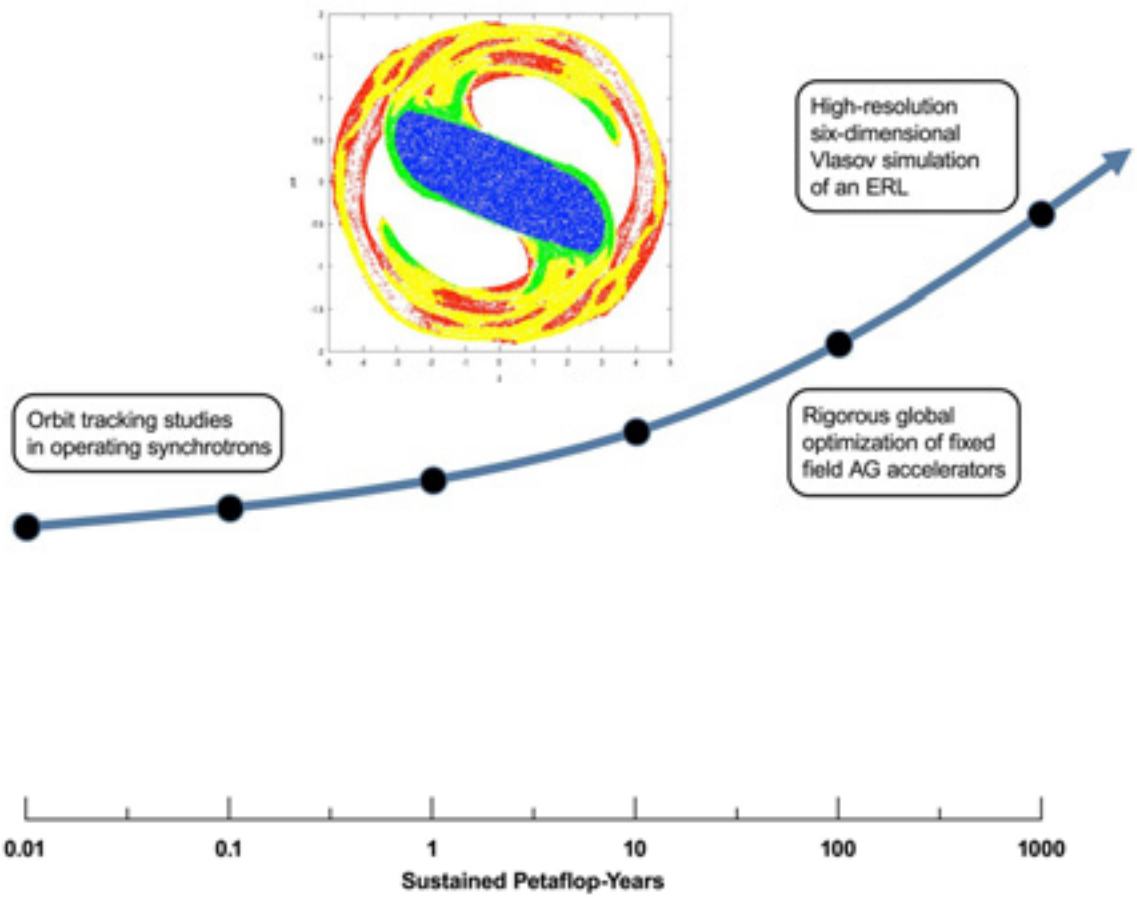

Figure 59. Anticipated highlights for priority research direction "Advanced Methods and Applications of Accelerator Simulations for Nuclear Physics Facilities." Image courtesy of Robert Ryne (Lawrence Berkeley National Laboratory).

The coherent electron-cooling concept proposes to combine the best features of electron cooling and stochastic cooling via free-electron laser technology. These features are combined to cool high-energy hadron beams on orders-of-magnitude shorter timescales (Litvinenko et al. 2008) than now possible. In a standard electron cooler, the key physical process is dynamic friction on the ions. The modulator section of a coherent cooler would be very similar to a standard cooler—but in this, case dynamical friction becomes irrelevant and the key physics is the shape of the density wake imprinted on the electron distribution by each ion. This implies use of a high-resolution PIC approach instead of the hybrid molecular dynamics/PIC approach used for conventional electron-cooling simulation. Though the high-resolution case involves approximately 1 billion cells, the time per step is comparable to the conventional case because resolutions of binary collisions are not required. As a result, the resource requirement is approximately the same as the conventional cooling scenario, namely approximately 10-petaflop-years over the course of 1 calendar year. The Poisson solver is a potential bottleneck, and it will be important that the solver for a single-point simulation scales to at least 10,000 computing cores.

Advanced Algorithms for Simulation and Optimization

The availability of extreme scale computing resources with massive memories will provide an opportunity to employ numerical algorithms that cannot realistically be implemented on current hardware. An example is the use of a direct Vlasov approach for the high-resolution prediction of ultra-low beam loss. Direct numerical simulation of the Vlasov equation has several advantages over the standard PIC approach. It allows for a direct and accurate description of particle distribution in the phase space; avoids random fluctuations and numerical noise caused by the finite number of macroparticles in a conventional PIC code; and enables scientists to accurately model low-density regions of the phase space such as the beam halo. Pure instability modes, which are difficult to study, are also conveniently simulated using PIC codes because the particle noise necessarily excites a spectrum of modes. Alternatively, the Vlasov approach is challenging computationally because calculations are performed in a (2n)-dimension phase space, where n = 1, 2, and 3; thus, the computational grid requires extreme computing resources. Current efforts have focused mainly on the two-dimensional (four-dimensional phase space) problem. To simulate multicomponent beams and plasma in a six-dimensional phase space, extreme scale computing will be necessary. For example, using 100 grid points in each direction requires 10^{12} grid points. The application of the Vlasov solver is an elegant approach to simulate three-dimensional problems in ion sources, be they electron cyclotron resonance or electron beam ion sources. These sources are inherent parts of the proposed FRIB.

Extreme scale computing resources will also provide the opportunity for employing new optimization algorithms. As an example, consider rigorous global optimization. Global parameter optimization will play a major role in the design of future nuclear physics particle accelerators. Recent advances in this rapidly progressing field allow the determination of rigorous solutions for global optimization problems with constraints. Based on dynamic domain decomposition and divide-and-conquer approaches, the methods are based on iteratively splitting the search space. On each of the currently active regions, a local or semilocal search is performed based on a variety of techniques including local descent-based method, linear or quadratic bounding, or genetic algorithms. More importantly, the currently active region is evaluated to determine whether, due to the behavior of the objective functions or the violation of constraints, it can be safely concluded not to contain the optimum, in which case it is discarded. While global parameter optimizations are very significant for the development of future particle accelerators, their computational requirements are immense. The underlying divide-and-conquer methods are currently performed on the largest available clusters and lend themselves to massive parallelization on 10^6 or more computing cores. One robust paradigm for parallelization is communication and load balancing on the "regular meeting" concept, at which all processors or suitable subgroups of processors share updates of the bounds of the objective function and redistribute large, unprocessed regions to achieve load balancing. Considering that a single evaluation of the objective function comprises a full simulation of the device in question—often comprising hundreds or thousands of core-hours—and that an exhaustive search can involve thousands to, in extreme cases, millions of evaluations of the objective function, tasks may result that require exaflop-months.

Advanced Tools for Accelerator Commissioning, Operation, and Control

Today, no accelerator facility in the world can fully rely on a computer model for its operation. There are intensive efforts to construct the needed modeling environment using experience gained at current state-of-the-art accelerators such as the Spallation Neutron Source, RHIC, and the Jefferson Laboratory. A promising approach is a configurable framework, independent of any specific accelerator that could allow

insertion of computing modules for different scale simulation, optimization, and control algorithms. Careful development of portable and mutable software types, or classes, to describe the accelerator elements is part of this and is another area for fruitful collaboration with applied mathematicians and computer scientists (Malitsky et al. 2009).[1] As more intense and complex machines are built, the challenge of achieving their design goals and operating them effectively grows larger. Using traditional operating and beam-tuning methods (mostly manual) will certainly result in commissioning delays and reduced machine availability for science. A complex accelerator system, as is being proposed for FRIB, uses primary beams ranging from proton beams to uranium beams, at different energies to produce secondary beams of rare isotopes spanning the whole chart of nuclides. Such an accelerator will require currently unavailable advanced tools to support operations. Using a parallel, realistic three-dimensional beam dynamics model to support commissioning and real-time machine operation will expedite the delivery of the machine, greatly enhance its availability, and significantly reduce its operating budget. Supporting online machine operations involves large-scale optimization problems; fast interfaces between the computer model, the beam diagnostic devices, and the beam line elements; and advanced data-analysis and -visualization tools. Combined with the requirement of fast turn-around calculations (seconds to minutes) to support fast decision making, the task can easily reach the extreme scale. This concept of a model-driven accelerator will benefit existing—and more importantly future—nuclear physics facilities.

Scientific Outcomes and Impacts

The development of advanced methods for extreme scale modeling of nuclear physics accelerators, and the application of these methods to novel types of accelerators, will have a major impact on future nuclear physics accelerator facilities. It will enable the exploration and possible usage of novel types of accelerators and accelerator systems (such as fixed-field alternating gradients, new methods of beam cooling, etc.). These methods will also lead to advanced design optimization techniques that will also benefit other types of DOE SC projects and lead to new tools for real-time control of large scientific experiments and facilities.

[1] Wang N. "Next Generation Communication Middleware for High Level Accelerator Control Systems." Paper presented at the *2009 International Computational Accelerator Physics Conference* in San Francisco, California. Not publicly available.

CONCLUSIONS AND RECOMMENDATIONS

CONCLUSIONS

Five major areas of nuclear physics have been discussed within this report: Cold Quantum Chromodynamics and Nuclear Forces, Nuclear Structure and Reactions, Nuclear Astrophysics, Hot and Dense Quantum Chromodynamics, and Accelerator Physics. This report makes clear that within each area, extreme scale computational resources are required to accomplish the objectives of the nuclear physics research program. Priority research directions that are crucial to the development of nuclear physics and that require extreme scale computing facilities have been identified in each of these areas. Collaborative efforts with experts in computer science, applied mathematics, statistics, high-energy physics, and other specialties outside nuclear physics have been determined to be crucial to fulfilling the mission of nuclear physics.

Nuclear physics is a field that has traditionally evolved from and flourished through a highly coordinated interplay between extensive experimental and theoretical programs. The overarching mission of the field is to establish a framework with which to perform high-precision calculations with quantifiable uncertainties of the properties and interactions of nuclear matter under a broad range of conditions, including those beyond the reach of laboratory experiment. Impressive experimental facilities, many of them employing accelerators to deliver beams of particles and nuclei, have provided many of the discoveries in the field. These facilities have helped researchers guide the theoretical developments that currently underpin the field. The next generation of accelerators is expected to provide a better understanding of nuclear physics, and to precisely measure nuclear properties and interactions that further constrain existing theoretical constructions. Since the 1970s, quantum chromodynamics (QCD) has been established as the theory of the strong interactions—which, along with the other forces of nature, is responsible for all nuclear phenomena. However, nuclear physics is a field defined by the regime of QCD in which its defining feature—asymptotic freedom—is concealed by confinement and by the vacuum; the numerical technique of lattice quantum chromodynamics (LQCD) is the only known way to perform *ab initio* QCD calculations of strong-interaction quantities in this regime. Remarkable progress has been made in the last decade in understanding the structure of hadrons and in the use of effective field theory to establish bridges between QCD, nuclear structure, and nuclear reactions. The 12 GeV upgrade at the Thomas Jefferson National Accelerator Facility is designed in part to enable the discovery and exploration of exotic hadrons for which the gluons of QCD play a visible role in their structure. The QCD-consistent forces that have emerged from the effective field theory framework are currently being used to calculate the properties and interactions of light nuclei. These forces are foreseen to be a central component of future calculations of all processes involving nucleons. Through significant collaborative efforts within the nuclear structure community (the Universal Nuclear Energy Density Functional collaboration), great strides are being made in developing nuclear many-body techniques that will enable reliable calculability throughout the periodic table of elements and of nuclear systems in extreme environments. Given that the nuclear systems range from the very simple to the very complex, a broad array of theoretical techniques are being developed. This continued development is crucial to fully use the results obtained with the experimental program and to make the connection with QCD. With these developments, the field of nuclear physics is presently entering an era in which precise QCD-based calculations of the properties and interactions of nuclei with quantifiable uncertainties will become possible. Extreme scale computing capabilities are required to perform the LQCD, the nuclear structure, and the nuclear reactions calculations that will achieve this revolutionary unification of nuclear physics.

Extreme scale computing is required both to verify (by comparison with experimental data) and to predict where there will be little or no experimental guidance.

Results that emerged from the experimental program at the Relativistic Heavy Ion Collider at Brookhaven National Laboratory and other heavy ion facilities led to dramatic progress in understanding the evolution of the hot and dense QCD matter from the initial collisions of heavy nuclei to the fragments that enter the detectors. A new state of matter with a low ratio of entropy-to-viscosity appears to have been produced during the earliest moments of the collision. The experimental program has stimulated huge advances in the theoretical description of such processes. The complexity of such a collision requires employing a range of theoretical techniques during the somewhat distinct phases of the evolution, including LQCD; nonequilibrium, non-Abelian plasma dynamics; and viscous hydrodynamics—all of which correspond to different regimes of QCD. A detailed map of the QCD phase-diagram—and reliable calculations with quantifiable uncertainties of the evolution of such systems through each of these phases and their subsequent experimental verification—will require extreme scale computing. Extreme scale capabilities are required to perform calculations of nuclear matter under the even more-extreme conditions of temperature and density for which experiments are not possible, but which existed at the very earliest times of our universe and may exist in other extreme environments.

Simulations of dynamic astrophysical systems, such as core-collapse supernovae, necessarily involve the inclusion of phenomena acting over an extreme range of length-scales from the nuclear to the stellar and including physics areas from nuclear reactions to large-scale turbulence in electromagnetic plasmas. Consequently, high-performance computing has always played a central role in astrophysical simulations. Complete simulations of such systems are required to be three-dimensional, including an accurate description of turbulence at all scales and dynamics dictated by local particle kinetics. Extreme scale computing resources are required to accurately describe the phenomena occurring over the large range of length-scales that are important in these simulations. QCD-based calculations—including those of nuclear material at finite temperatures and densities, as well as of nuclear structure and reactions that become possible only with extreme scale computing capabilities—will improve the reliability of the input to these astrophysical simulations, enabling quantification of uncertainties in such simulations. While providing a profound level of understanding of astrophysical phenomena including neutron-star mergers and X-ray bursters, astrophysical simulation at the extreme scale will constrain the properties and interactions of neutrinos and improve overall understanding of the very earliest times in the universe.

Extreme scale computing will enable predictive capability for the theory of fission. This will give scientists the ability to calculate the fission cross-sections of major (e.g., ^{235}U, ^{238}U, and ^{239}Pu) and minor actinides and their by-products—the latter two of which are poorly known from experiment because of their short half-lives. This knowledge is important for the design of our nation's next generation of fission (breeder) reactors, as it will allow for their energy production, waste production, and efficiency to be quantified with reliable uncertainties. A clearer understanding of the waste produced by these reactors will reduce the uncertainty in the amount of unspent nuclear fuel that remains in the reactors as time passes which, in turn, will guide nonproliferation policies that ensure nuclear material within these reactors remains secure. This information is an important element in addressing national energy concerns.

With extreme scale computing revolutionizing nuclear physics through the unification of efforts that are presently operationally disconnected, a deeper cross-fertilization between these areas will lead to further scientific breakthroughs as a result of deeper understanding gained from the results of the extreme scale

computations. High-performance computing will become as important to the nuclear physics program as the experimental and theoretical components. Just as progress in experiment and theory are tightly connected, so will be the case for extreme scale computing, nuclear theory, and experiments. Perhaps more importantly, extreme scale computing resources will generate enhancements to the nuclear physics program that cannot be imagined today. Resulting from extreme scale computing, breakthroughs will occur that will allow for significantly more progress to be made than can be presently estimated—including development of new areas of nuclear physics research. Though unable to predict the discoveries that will occur from a tighter integration of extreme scale computing with the nuclear physics community, researchers can state that quite novel technologies will be developed that will augment nuclear physics capabilities. Historically, theoretical and experimental nuclear physics have played complementary roles in advancing the field of nuclear physics. With the inclusion of extreme scale computing, the rate of advancement will increase. This is the realization of computational research as a third leg of scientific exploration within nuclear physics.

Nuclear physics at the extreme scale will also require an array of smaller computational facilities, with an aggregate capacity comparable to what will be used at the extreme scale. As is the case today, development and optimization of codes and algorithms will be on smaller machines that are located at either national computational facilities or that are local to researchers at universities or laboratories. Post-processing the results generated at the extreme scale facilities will, for many research areas, require a broad range of capabilities and capacities at 1% to 10% of the extreme scale. This is also true of the pre-processing that will occur in preparation for running on the extreme scale facilities.

More researchers will be needed to undertake the physics-, algorithm-, and code-developments that must be accomplished before the deployment of extreme scale computing, currently estimated to be around 2017. As the evolution to the extreme scale is expected to see computing resources grow faster than Moore's Law predicts, enhanced recruitment and education mechanisms must be identified and implemented. Consensus estimates originating from the subpanels suggest that, in addition to present-day personnel, approximately 10 graduate students and 10 postdoctoral fellows will be required for each of the identified extreme scale projects.

While extreme-scale computing will provide the computational resources required to perform calculations that will unify nuclear physics research, the collaborations between the subfields of nuclear physics must be substantially enhanced to practically implement this unification. Extreme-scale computing will be a catalyst for such collaborations. If completed appropriately, these collaborative efforts will extend beyond the traditional borders of nuclear physics to include particle physicists, plasma physicists, fluid dynamicists, and other researchers. The nuclear physics community has venues from which to launch such collaborative efforts. These venues support short workshops and, in the case of the Institute for Nuclear Theory, longer programs that have the potential for launching new areas of research or that bring together scientists from different areas of research in a collaborative forum to generate progress in one or more of these areas. Such collaboration and innovation are required for the nuclear physics community to prepare for, and make optimal use of, extreme scale computational resources.

Nuclear physicists are benefiting greatly from interactions and collaborations with computer scientists, and such collaborations must be further embraced and strengthened as researchers move toward the era of extreme computing. Given the somewhat different hardware requirements between, for example, nuclear structure and reactions, and LQCD, it is imperative that nuclear physicists work closely with the computer scientists involved in designing the extreme-scale computing environments and all aspects of future

hardware and software development. Further, it is imperative that nuclear physicists be provided access to machine design and machine simulators to ensure codes are ready for production when new machines become available. Given the sheer magnitude of extreme scale computational machines, present-day codes will not be viable, and collaborative efforts with computer scientists will be essential in developing new coding paradigms for the extreme computing era. Major progress has occurred in large-scale numerical calculations in nuclear physics through the implementation of procedures and techniques that have been developed by, or in collaboration with, applied mathematicians. Collaborations continue to be important to algorithm and code development in nuclear physics, and it is anticipated that collaborations with statisticians are likely to become a vital component of the extraction of physics from the large data output of extreme scale computations.

Future investments in dedicated centers that house researchers in multidisciplinary fields that use high-performance computers will be essential for performing research on the next generation of computers. Applied mathematicians, computational scientists, and nuclear physicists must work together to design, port, and implement codes on next-generation computers. Hardware and algorithmic issues that are created from the sheer size and scope related to extreme scale computing can only be tackled through multidisciplinary teams. In many ways, our nation's national laboratories can play an integral role in promoting nuclear physics research via high-performance computing because they have the infrastructure and requisite multidisciplinary fields in place. As the extreme era evolves, however, centers dedicated to extreme scale computing must be developed and institutionalized. Ideally, these centers should draw upon resources already in place at national laboratories.

Many of the issues discussed in this section have, in no small way, motivated the Scientific Discovery through Advanced Computing (SciDAC) initiative. This panel uniformly views SciDAC as a great success and a central component of present and future research in nuclear physics. Preparation for extreme scale computing will require nurturing and strengthening the collaborations built via the SciDAC programs, and the knowledge gained from implementing SciDAC will be useful in planning for the extreme scale era. It is clear, however, that the SciDAC program alone will not be sufficient to bring about the extreme scale era. A program is needed that has similar goals but has broader size and scope. This program must also be stable on a timescale set by the full maturation of extreme scale computing (~10 years or more). An enhanced support mechanism will be required to establish and maintain the infrastructure of such a program and to further the multidisciplinary collaboration that will ultimately accomplish the mission of nuclear physics. As stated above, this program will extend beyond traditional borders and should be multidisciplinary and international in nature.

Extreme scale computing will lead to extreme scale data management issues. The post-computation data outputs and data manipulations will not resemble those that occur today. It is clear that less "data per flop" can be generated and stored. As a result, a great deal of sophistication must be developed in the nuclear physics community, in collaboration with computer scientists, to determine appropriate outputs. Even with such enhancements, the sheer volume of output will require new techniques for scientific discovery. Visualization is playing an important role in processing the output of large-scale calculations in many areas, but is not used as much in nuclear physics. With the help of computer scientists, this will change as the extreme scale era unfolds.

Many of the most important frontiers in nuclear physics research require extreme scale computational facilities to accomplish proposed objectives. Remarkably, extreme scale computing will enable an era of nuclear physics where theoretical efforts in nuclear physics are unified. The intrinsic structure and

interactions of the lightest hadrons, including the nucleons, will be unambiguously determined. It will become practical to perform calculations underpinned by QCD, in areas of research such as astrophysics, nuclear structure and reactions, and in heavy ion collisions. These calculations will address central issues in nuclear physics, such as the evolution of a heavy ion collision from the time of collision until detection; the fine-tuning in the triple-alpha reaction that produces ^{12}C and how the reaction depends upon the fundamental constants of nature; the precise calculation of very low-energy fusion cross-sections of light-nuclei; and the complete three-dimensional simulation of core-collapse supernovae.

Extreme scale computing will bring nuclear physics into an era where predictive capability will be typical. This era will bring with it answers to longstanding nuclear physics questions that not only impact key issues in fundamental science, but those that will help address our nation's security and energy needs. Investment in extreme scale computing resources will enable the United States to maintain its forefront position in nuclear science research.

RECOMMENDATIONS

A number of recommendations emerged from each of the subpanels. These recommendations were combined where possible into the following general set of recommendations.

Develop and Deploy Extreme Scale Computational Resources

- *Develop and deploy extreme scale computing facilities dedicated to the mission of nuclear physics.* As detailed in this report, the nuclear physics research program requires extreme scale computing to accomplish its mission. Extreme scale computing facilities should be designed to accommodate the needs of nuclear physicists and should be deployed to focus on the mission of nuclear physics. Associated with such computational resources, firm support to research and the development of technologies—including applied mathematics, computer science, parallel algorithms, parallel optimization, parallel numerical libraries, supporting software, data storage, data analysis, visualization and visual analytics—is also recommended.

- *Provide nuclear physicists early access to the design and development of extreme scale computing hardware.* A transparent process should be established to facilitate decisions regarding the nature of extreme scale computing. Active dialogue between computer system designers and the nuclear physics community is highly encouraged to optimize the nuclear physics output of the extreme scale facilities. An open format in the decision-making process for extreme scale machines, and those leading to the extreme scale, will allow sufficient time for codes to be adapted and tuned to the new machines, and for new ideas to be explored.

Access (Early, Range, Enhanced)

- *Provide nuclear physicists access to a range of computational resources, including local resources, during and after the development of extreme scale computing*. Supporting computational facilities with a range of capabilities and capacities is required to maximize the scientific output of extreme scale computing facilities. It is not computationally efficient, and in some cases, not computationally feasible to apply all algorithms on a single extreme scale computer. Therefore, a single extreme scale computer will be limited in scientific output without additional significant computational resources that will be used for pre- and post-processing of numerical "data" as well as algorithm development

and optimization. The range of capabilities of these smaller computers could potentially span several orders of magnitude while still representing a large fraction of the total computational capacity.

- *Provide nuclear physicists enhanced access to computing resources and to the most powerful computers from this point forward.* Preparation for extreme scale computing requires a dedicated effort in the development and optimization of algorithms and codes. This will require continuous access to computational resources that are significantly larger than those presently available to the nuclear physics community at major computing facilities. Even with the allocations the scientific community has received under the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, nuclear physicists are currently not obtaining enough computer time to accomplish its present-day objectives. It is strongly recommended that computing resources available to the nuclear physics community be substantially enhanced (much greater than an order of magnitude) from this point forward. Further, the present mode of operation that requires collaborations to prepare and submit multiple proposals for computer resources during any given year is suboptimal. Nuclear physics requires continuous access to substantial computing resources.

Collaboration, Training, and the Next Generation

- *Foster collaboration between nuclear physicists, computational scientists, applied mathematicians, and physicists outside of nuclear physics.* The efficient use of extreme scale machines will require a diverse set of talents including understanding physics applications, algorithmic invention, programming skills, and understanding the new computer systems' hardware and architecture. Communities involving nuclear physics, computer science, and applied math must collaborate to develop the necessary hardware, algorithms, and software infrastructure. A renewed focus on investment and collaboration between different scientific disciplines will be crucial in assembling the human resources required for the extreme computing era.
- *Strengthen the infrastructure for support and training at a range of career levels.* Optimal use of extreme scale computing facilities will require training of its potential users to begin in the near term. This training will be required at various stages throughout a user's career. Graduate-level courses and degrees that further educate nuclear physicists in key issues in computer science and applied mathematics are envisioned, while the training of researchers in nuclear physics and high-performance computing at the post-graduate level will occur at summer schools. For more senior researchers, such as faculty members at universities or staff members at national laboratories, such infrastructure is already in place, as programs and workshops at the Institute for Nuclear Theory (and possibly elsewhere) can be initiated for training or retraining purposes. An increase in funding for research associates and for related SciDAC programs is expected to be required to meet the challenges in software development and data handling that will accompany the extreme scale era.
- *Provide long-term positions to researchers to encourage interdisciplinary research that will enhance the nuclear physics programs at the extreme scale.* It is important to establish career paths for new researchers in interdisciplinary areas related to extreme scale computing so that their efforts in solving the challenges of extreme computing do not limit their future careers. These researchers could be located at national laboratories or at computational-science centers. Joint laboratory/university positions could address many concerns of a new researcher. Joint positions would also produce an open recruitment channel between universities and laboratories. Such positions would give laboratory staff access to students, and universities access to unique laboratory resources.

REFERENCES

Aarts G and IO Stamatescu. 2008. "Stochastic Quantization at Finite Chemical Potential." *Journal of High Energy Physics* 0809:018 DOI: 10.1088/1126-6708/2008/09/018.

Advanced Fuel Cycles (AFC). 2006. *Report of the Nuclear Physics and Related Computational Science R&D for the Advanced Fuel Cycles Workshop*, August 10-12, Bethesda, Maryland. Accessed May 10, 2009, at http://www.mcs.anl.gov/events/workshops/nprcsafc/Report_FINAL.pdf.

Alexakis A, AC Calder, A Heger, EF Brown, LJ Dursi, TW Truran, R Rosner, DQ Lamb, FX Timmes, B Fryxell, M Zingale, PM Ricker, and K Olson. 2004. "On Heavy Element Enrichment in Classical Novae." *The Astrophysical Journal* 602(2):931-937.

Alexandrou C, Ph de Forcrand, Th Lippert, H Neff, JW Negele, K Schilling, W Schroers, and A Tsapalis. 2004. "N to D Electromagnetic Transition Form Factors from Lattice QCD." *Physical Review D* 69(11): Art. No. 114506. DOI:10.1103/PhysRevD.69.114506. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.69.114506.

Alexandrou C, Ph de Forcrand, H Neff, W Negele, W Schroers, and A Tsapalis. 2005. "The N to Delta Electromagnetic Transition Form Factors from Lattice QCD." *Physical Review Letters* 94(2): Art. No. 021601. DOI:10.1103/PhysRevLett.94.021601. Accessed May 10, 2009, at http://arxiv.org/abs/hep-lat/0409122.

Alexandrou C, G Koutsou, H Neff, JW Negele, W Schroers, and A Tsapalis. 2008. "The Nucleon to Delta Electromagnetic Transition Form Factors in Lattice QCD." *Physical Review D* 77(8): Art. No. 085012.

Alexandru A, M Faber, I Horvath, and K-F Liu. 2005. "Lattice QCD at Finite Density via a New Canonical Approach." *Physical Review D* 72(11): Art. No. 114513. DOI:10.1103/PhysRevD.72.114513. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.72.114513.

Alford MG, A Schmitt, K Rajagopal, and T Schafer. 2008. "Color Superconductivity in Dense Quark Matter." *Reviews of Modern Physics* 80:1455-1515. DOI:10.1103/RevModPhys.80.1455. Accessed May 10, at http://arxiv.org/abs/hep-ph/0003185.

Allton CR, S Ejiri, SJ Hands, O Kaczmarek, F Karsch, E Laermann, and C Schmidt. 2003. "The Equation of State for Two Flavor QCD at Non-Zero Chemical Potential." *Physical Review D* 68:014507. Accessed May 10, 2009, at http://arxiv.org/abs/hep-lat/0305007.

Almgren AS, JB Bell, CA Rendleman, and M Zingale. 2006. "Low Mach Number Modeling of Type Ia Supernovae. I. Hydrodynamics." *The Astrophysical Journal* 637(2):922-936.

Anderson HC. 1980. "Molecular Dynamics Simulations at Constant Pressure and/or Temperature." *The Journal of Chemical Physics* 72: Art. No. 2384. DOI:10.1063/1.439486.

Angulo C, M Arnould, M Rayet, P Descouvemont, D Baye, C Leclercq-Willain, A Coc, S Barhoumi, P Aguer, C Rolfs, R Kunz, JW Hammer, A Mayer, T Paradellis, S Kossionides, C Chronidou, K Spyrou, S Degl'Innocenti, G Fiorentini, B Ricci, S Zavatarelli, C Providencia, H Wolters, J Soares, C Grama, J Rahighi, A Shotter, and ML Rachti. 1999. "A Compilation of Charged-Particle Induced Thermonuclear Reaction Rates." *Nuclear Physics A* 656(1):3-183. DOI: 10.1016/S0375-9474(99)00030-5.

Arnett D. 1996. *Supernovae and Nucleosynthesis. An Investigation of the History of Matter, from the Big Bang to the Present.* Princeton Series in Astrophysics, Princeton University Press, Princeton, New Jersey.

Arnett WD. 1966. "Gravitational Collapse and Weak Interactions." *Canadian Journal of Physics* 44(11):2553-2594.

Arnett WD. 1967. "Mass Dependence in Gravitational Collapse of Stellar Cores." *Canadian Journal of Physics* 45(5):1621-1641.

Arnett WD. 1969. "A Possible Model of Supernovae: Detonation of ^{12}C." *Astrophysics and Space Science* 5(2):180-212.

Arnold P, GD Moore, LG Yaffe. 2003. "Effective Kinetic Theory for High Temperature Gauge Theories." JHEP 0301:030, 2003. e-Print: hep-ph/0209353.

Arnold P, J Lenaghan, GD Moore, LG Yaffe. 2004. "Apparent Thermalization Due to Plasma Instabilities in the Quark-Gluon Plasma." *Physical Review Letters* 94:072302, 2005. e-Print: nucl-th/0409068.

Arve P, GF Bertsch, JW Negele, and G Puddu. 1987. "Model for Tunneling in Many-Particle Systems." *Physical Review C* 36(5):2018-2025. DOI:10.1103/PhysRevC.36.2018. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.36.2018.

Asakawa M, T Hatsuda, and Y Nakahara. 2001. "Maximum Entropy Analysis of the Spectral Functions in Lattice QCD." *Progress in Particle Nuclear Physics* 46(2):459-508. DOI:10.1016/S0146-6410(01)00150-8.

Assuncão M, M Fey, A Lefebvre-Schuhl, J Kiener, V Tatischeff, JW Hammer, C Beck, C Boukari-Pelissie, A Coc, JJ Correia, S Courtin, F Fluerot, E Galanopoulos, C Grama, F Haas, F Hammache, F Hannachi, S Harissopulos, A Korichi, R Kunz, D LeDu, A Lopez-Martens, D Malcherek, R Meunier, T Paradellis, M Rousseau, N Rowley, G Staudt, S Szilner, JP Thibaud, and JL Weil. 2006. "$E1$ and $E2$ S Factors of $^{12}C(\alpha, \gamma_0)^{16}O$ from γ-Ray Angular Distributions with a 4 π-Detector Array." *Physical Review C* 73(5):055801-1–055801-19.

Avignone FT, SR Elliot, and J Engel. 2008. "Double Beta Decay, Majorana Neutrinos, and Neutrino Mass." Reviews of Modern Physics 80(2): Art. No. 481. DOI:10.1103/RevModPhys.80.481. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/RevModPhys.80.481.

Bahcall JN, MH Pinsonneault, and S Basu. 2001. "Solar Models: Current Epoch and Time Dependences, Neutrinos, and Helioseismological Properties." *The Astrophysical Journal* 555(2, Part 1):990-1012.

Bailey DH. 2008. "High Precision Numberical Integration: Progress and Challenges." *Proceedings of the Milestones in Computer Algebra*. LBNL-547E. Lawrence Berkeley National Laboratory, California.

Baruah S, G Audi, K Blaum, M Dworschak, S George, C Guénaut, U Hager, F Herfurth, A Herlert, A Kellerbauer, H-J Kluge, D Lunney, H Schatz, L Schweikhard, and C Yazidjian. 2008. "Mass Measurements Beyond the Major *r*-Process Waiting Point ^{80}Zn." *Physical Review Letters* 101(26):262501-1–262501-4.

Basak S, R Edwards, GT Fleming, UM Heller, C Morningstar, D Richards, I Sato, and SJ Wallace. 2005a. "Clebsch-Gordan Construction of Lattice Interpolating Fields for Excited Baryons." *Physical Review* D72(7): Art. No. 074501. DOI:10.1103/PhysRevD.72.074501. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.72.074501.

Basak S, RG Edwards, GT Fleming, UM Heller, C Morningstar, D Richards, I Sato, and S Wallace. 2005b. "Group-Theoretical Construction of Extended Baryon Operators in Lattice QCD." *Physical Review* D72(9):094506.1-094506.15.

Bass SA, B Mueller and W Poeschl. 1999. "Lattice Gauge Description of Colliding Nuclei." *Journal of Physics G: Nuclear and Particle Physics* 25 L109-L115. DOI:10.1088/0954-3899/25/11/101.

Bazavov A, T Bhattacharya, M Cheng, NH Christ, C DeTar, S Ejiri, S Gottlieb, R Gupta, UM Heller, K Huebner, C Jung, F Karsch, E Laermann, L Levkova, C Miao, RD Mawhinney, P Petreczky, C Schmidt, RA Soltz, W Soeldner, R Sugar, D Toussaint, and P Vranas. 2009. "Equation of State and QCD Transition at Finite Temperature." *High Energy Physics – Lattice*. Accessed May 10, 2009, at http://arxiv.org/abs/0903.4379v1.

Beane SR, P Bedaque, K Orginos, and M Savage. 2006a. (NPLQCD Collaboration). "Nucleon-Nucleon Scattering from Fully-Dynamical Lattice QCD." *Physical Review Letters* 97(1):012001.

Beane SR, P Bedaque, T Luu, K Orginos, E Pallente, A Parreno, and M Savage. 2006b. (NPLQCD Collaboration). "π K Scattering in Full QCD with Domain-Wall Valence Quarks." *Physical Review D* 74: Art. No. 114503. DOI:10.1103/PhysRevD.74.114503.

Beane SR, PF Bedaque, TC Luu, K Orginos, E Pallante, A Parreno, and MJ Savage. 2007. (NPLQCD Collaboration). "Hyperon-Nucleon Scattering from Fully-Dynamical Lattice QCD." *Nuclear Physics A* 794:62-72. DOI:10.1016/j.nuclphysa.2007.07.006.

Beane SR, TC Luu, K Orginos, A Parreno, MJ Savage, A Torok, and A Walker-Loud. 2008a. (NPLQCD Collaboration). "Precise Determination of the I=2 $\pi\pi$ Scattering Length from Mixed-Action Lattice QCD." *Physical Review D* 77(1): Art. No. 014505. DOI:10.1103/PhysRevD.77.014505.

Beane SR, W Detmold, TC Luu, K Orginos, MJ Savage, and A Torok. 2008b. (NPLQCD Collaboration). "Multi-Pion Systems in Lattice QCD and the Three-Pion Interaction." *Physical Review Letters* 100(8): Art. No. 082204. DOI:10.1103/PhysRevLett.100.082004. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.100.082004.

Beane SR, K Orginos, and MJ Savage. 2008c. "Hadronic Interactions from Lattice QCD." *International Journal of Modern Physics E* 17(7):1157-1218. DOI:10.1142/S0218301308010404.

Bell G, D Bruhwiler, A Fedotov, A Sobol, R Busby, P Stoltz, D Abell, P Messmer, I Ben-Zvi, and V Litvinenko. 2008. "Simulating the Dynamical Friction Force on Ions due to a Briefly Co-Propagating Electron Beam." 2008. *Journal of Computational Physics* 227:8714-8735.

Bertsch GF, DJ Dean, and W Nazarewicz. 2007. "Computing Atomic Nuclei." *SciDAC Review, Winter 2007* 6:42-51. Accessed May 10, 2009, at http://www.scidacreview.org/0704/html/unedf.html.

Berz M. 1999. *Modern Map Methods in Particle Beam Physics*. Academic Press, San Diego, California. Accessed May 10, 2009, at http://bt.pa.msu.edu/pub/.

Berz M, K Makino, and Y-K Kim. 2005. "Long-Term Stability of the Tevatron by Validated Global Optimization." *Nuclear Instruments and Methods* 558(1):1-10.

Binder K. 1981. "Critical Properties from Monte Carlo Coarse Graining and Renormalization." *Physical Review Letters* 47(9):693-696. DOI:10.1103/PhysRevLett.47.693. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.47.693.

Blondin JM and A Mezzacappa. 2006. "The Spherical Accretion Shock Instability in the Linear Regime." *The Astrophysical Journal* 642(1):401-409.

Borasoy B, E Epelbaume, H Krebs, D Lee, and U-G Meissner. 2007. "Lattice Simulations for Light Nuclei: Chiral Effective Field Theory at Leading Order." *The European Physical Journal A, Hadrons and Nuclei* 31(1):105-123.

Borasoy B, E Epelbaume, H Krebs, D Lee, and U-G Meissner. 2008. "Dilute Neutron Matter on the Lattice at Next-to-Leading Order in Chiral Effective Field Theory." *The European Physical Journal A, Hadrons and Nuclei 35(3):*357-367.

Branscomb L, T Belytschko, P Bridenbaugh, T Chay, J Dozier, GS Grest, EF Hayes, B Honig, N Lane, WA Lester Jr, GJ McRae, JA Sethian, B Smith and M Vernon. 1993. *NSF Blue Ribbon Panel on High Performance Computing:* "From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing." NSB 93-205. National Science Foundation, Washington, D.C. Accessed September 1, 2009, at http://www.nsf.gov/pubs/stis1993/nsb93205/nsb93205.txt.

Brommel D, M Göckeler , Ph Hägler , R Horsley, Y Nakamura, M Ohtani, D Pleiter, PEL Rakow, A Schäfer, G Schierholz, W Schroers, H Stuben, and JM Zanotti. 2007. (QCDSF/UKQCD Collaborations). "Moments of Generalized Parton Distributions and Quark Angular Momentum of the Nucleon." The 25th International Symposium on Lattice Field Theory, *Proceedings of Science (Lattice 2007)*:158.

Brown GE, A Heger, N Langer, C-H Lee, S Wellstein, and HA Bethe. 2001. "Formation of High Mass X-ray Black Hole Binaries." *New Astronomy* 6 (7): 457-470 (2001).

Brun AS and J Toomre. 2002. "Turbulent Convection Under the Influence of Rotation: Sustaining a Strong Differential Rotation." *The Astrophysical Journal* 570(2):865-885.

Bulava JM, RG Edwards, E Engelson, J Foley, B Joó, A Lichtl, H-W Lin, N Mathur, C Morningstar, DG Richards, and SJ Wallace. 2009. "Excited State Nucleon Spectrum with Two Flavors of Dynamical Fermions." *Physical Review* D 79(3): Art. No. 034505. DOI:10.1103/PhysRevD.79.034505.

Burrows A, J Hayes, and BA Fryxell. 1995. "On the Nature of Core-Collapse Supernova Explosions." *The Astrophysical Journal* 45092(2, Part 1):830-850.

Calder AC, DM Townsley, IR Seitenzahl, F Peng, OEB Messer, N Vladimirova, EF Brown, JW Truran, and DQ Lamb. 2007. "Capturing the Fire: Flame Energetics and Neutronization for Type Ia Supernova Simulations." *The Astrophysical Journal* 656(1)313-332.

Cameron AGW. 1955. "Origin of Anomalous Abundances of the Elements in Giant Stars." *Astrophysical Journal* 121(1):144-60.

Camilleri L, E Lisi, and JF Wilkerson. 2008. "Neutrino Masses and Mixings: Status and Prospects." *Annual Review of Nuclear and Particle Science* 58(1):343-369. DOI:10.1146/annurev.nucl.57.090506.123038.

Chandrasekharan S and U-J Wiese. 1999. "Meron-Cluster Solution of a Fermion Sign Problem." *Physical Review Letters* 83(16):3116-3119. DOI: 10.1103/PhysRevLett.83.3116. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.83.3116.

Chang SY, J Morales Jr, VR Pandharipande, DG Ravenhall, J Carlson, SC Pieper, RB Wiringa, and KE Schmidt. 2004. "Neutron Matter: A Superfluid Gas." *Nuclear Physics A* 746:215. DOI: 10.1016/j.nuclphysa.2004.09.119.

Chen JW, DO Connell, and A Walker-Loud. 2007. "Two Meson Systems with Ginsparg-Wilson Valence Quarks." *Physical Review D* 75:054501. DOI:10.1103/PhysRevD.75.054501.

Chen P. 2001. "Heavy Quarks on Anisotropic Lattices: The Charmonium Spectrum." *Physical Review D* 64(3): Art. No. 034509. DOI:10.1103/PhysRevD.64.034509. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.64.034509.

Chorin AJ. 2008. "Monte Carlo Without Chains." *Communications in Applied Mathematics and Computational Science* 3: 77-93.

Colgate SA and RH White. 1966. "Hydrodynamic Behavior of Supernovae Explosions." *The Astrophysical Journal* 143(3):626-681.

Crawford CH, P Henning, M Kistler, and C Wright. 2008. "Accelerating Computing with the Cell Broadband Engine Processor." In *Proceedings of the 2008 Conference on Computing Frontiers*, pp. 3-9, May 5-7, 2008, Ischia, Italy.

D'Elia M and MP Lombardo. 2003. "Finite Density QCD via Imaginary Chemical Potential." *Physical Review D* 67(1): Art. No. 014505. DOI:10.1103/PhysRevD.67.014505. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.67.014505.

Davis R Jr. 2003. "Nobel Lecture: A Half-Century with Solar Neutrinos." *Reviews of Modern Physics* 75(3):985-994.

de Forcrand P and O Philipsen. 2002. "The QCD Phase Diagram for Small Densities from Imaginary Chemical Potential." *Nuclear Physics B* 642(1-2):290-306. DOI:10.1016/S0550-3213(02)00626-0. Accessed May 10, 2009, at http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6TVC-46K8PWG-1&_user=2741876&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000058656&_version=1&_urlVersion=0&_userid=2741876&md5=53396a285404ee887d10bceba3ffa226.

de Forcrand P. 2009. "The Curvature of the Critical Surface (m(u,d),m(s))**crit(mu): A Progress Report. (Zurich, ETH & CERN), Owe Philipsen (Munster U., ITP)." Presentation given at 26th International Symposium on Lattice Field Theory, Williamsburg, Virginia, July 14-20, 2008. Published in PoS LATTICE2008:208, 2008. e-Print: arXiv:0811.3858.

DeTar C. 2008. "Recent Progress in Lattice QCD Thermodynamics." In the *26th International Symposium on Lattice Field Theory (Lattice 2008)*, July 14-19, Williamsburg, Virginia.

Detmold W and MJ Savage. 2009. "Color Screening by Pions." *Physical Review Letters* 102(3): Art. No. 032004. DOI:10.1103/PhysRevLett.102.032004. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.102.032004.

Detmold W, MJ Savage, A Torok, SR Beane, TC Luu, K Orginos, and A Parreno. 2008a. (NPLQCD Collaboration) "Multi-Pion States in Lattice QCD and the Charged-Pion Condensate." *Physical Review D* 78:Art. No. 014507. DOI:10.1103/PhysRevD.78.014507. Accessed May 10, 2009, at http://arxiv.org/abs/0803.2728.

Detmold W, K Orginos, MJ Savage, and AWalker-Loud. 2008b. "Kaon Condensation with Lattice QCD." *Physical Review D* 78(5): Art. No. 054514. DOI:10.1103/PhysRevD.78.054514. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.78.054514.

DOE. 2003. *Facilities for the Future of Science: A Twenty-Year Outlook*. DOE/SC-0078. U.S. Department of Energy, Office of Science, Washington, D.C. Accessed September 1, 2009, at http://www.er.doe.gov/Scientific_User_Facilities/History/20-Year-Outlook-screen.pdf.

DOE. 2007. *The Frontiers of Nuclear Science, A Long Range Plan*. Prepared by the Nuclear Science Advisory Committee for the U.S. Department of Energy's Office of Science for Nuclear Physics and the National Science Foundation, Washington, D.C. Accessed June 4, 2009, at http://www.sc.doe.gov/np/nsac/docs/Nuclear-Science.High-Res.pdf.

DOE. 2008. *Report to the Nuclear Science Advisory Committee*. Prepared by the Subcommittee on Performance Measures/Nuclear Science Advisory Committee for the U.S. Department of Energy's Office of Science and the National Science Foundation, Washington, D.C. Accessed May 10, 2009, at http://www.sc.doe.gov/np/nsac/docs/PerfMeasEvalFinal.pdf.

Dudek JJ and RG Edwards. 2006. "Two Photon Decays of Charmonia from Lattice QCD." *Physical Review Letters* 97:172001.

Dudek JJ, RG Edwards, and DG Richards. 2006. "Radiative Transitions in Charmonium from Lattice QCD." *Physical Review D* 73: Art. No. 074507.

Dudek JJ, RG Edwards, N Mathur, and DG Richards. 2008. "Charmonium Excited State Spectrum in Lattice QCD." *Physical Review D* 77(3): Art. No. 034501. DOI:10.1103/PhysRevD.77.034501.

Dumitru A, Y Nara, and M Strickland. 2007. "Ultraviolet Avalanche in Anisotropic Non-Abelian Plasmas." *Physical Review D* 75(2):025016. DOI:10.1103/PhysRevD.75.025016. Accessed May 10, 2009, at http://arxiv.org/abs/hep-ph/0604149.

Edwards RG, GT Fleming, Ph Hägler, JW Negele, K Orginos, AV Pochinsky, DB Renner, DG Richards, and W Schroers. 2006a. (LHPC Collaboration). "The Nucleon Axial Charge in Full Lattice QCD." *Physical Review Letters* 96(5): Art. No. 052001.

Edwards RG, G Fleming, Ph Hägler, JW Negele, K Orginos, AV Pochinsky, DB Renner, DG Richards, and W Schroers. 2006b. (LHPC Collaboration). "Nucleon Structure in the Chiral Regime with Domain Wall Fermions on an Improved Staggered Sea." *Proceedings of Science (Lattice 2006)*:121.

Eggleton PP, DSP Dearborn, and JC Lattanzio. 2006. "Deep Mixing of ^{3}He: Reconciling Big Bang and Stellar Nucleosynthesis." *Science* 314(5805):1580-1583.

Ehmann C and G Bali. 2007. "Charmonium Spectrum Including Higher Spin and Exotic States." *Proceedings of Science (Lattice 2007)*: Art. No. 094, *International Symposium on Field Theory on the Lattice*, Amsterdam, The Netherlands.

Endrodi G, Z Fodor, SD Katz, and KK Szabo. 2007. "The Equation of State at High Temperatures from Lattice QCD." PoS (LAT2007) 228.

Essence. 2007. "ESSENCE: Measuring the Equation of State of the Universe." Accessed May 10, 2009, at http://www.ctio.noao.edu/~wsne/.

European Gravitational Observatory (EGO). 2006. *A Gravitational Waves Antenna: Virgo*. Accessed May 10, 2009, at http://www.ego-gw.it/virgodescription/.

Fedotov AV, DL Bruhwiler, AO Sidorin, DT Abell, I Ben-Zvi, R Busby, JR Cary, and VN Litvinenko. 2006. "Numerical Study of the Magnetized Friction Force." *Physical Review Special Topics−Accelerators and Beams* 9(7): Art. No. 074401. DOI:10.1103/PhysRevSTAB.9.074401. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevSTAB.9.074401.

Ferrenberg AM and RH Swendsen. 1988. "New Monte Carlo Technique for Studying Phase Transitions." *Physical Review Letters* 61(23):2635-2638. DOI:10.1103/PhysRevLett.61.2635. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.61.2635.

Follana E, Q Mason, C Davies, K Hornbostel, GP Lepage, J Shigemitsu, H Trottier, and K Wong. 2007. (HPQCD Collaboration and UKQCD Collaboration). "Highly Improved Staggered Quarks on the Lattice, with Applications to Charm Physics." *Physical Review D* 75:054502. DOI:10.1103/PhysRevD.75.054502. Accessed May 10, 2009, at http://arxiv.org/abs/hep-lat/0610092.

Frey L, CL Fryer, and AL Hungerford. 2009. "Simulations of Supernova Shock Breakout." *Bulletin of the American Astronomical Society*, Vol 41. The 213th Meeting of the American Astronomical Society with HAD and HEAD, January 4-8, 2009, Long Beach, California.

Frieman JA, MS Turner, and D Huterer. 2008. "Dark Energy and the Accelerating Universe." *Annual Review of Astronomy and Astrophysics* 46:385-432.

Fröhlich C, P Hauser, M Liebendörfer, G Martínez-Pinedo, F-K Thielemann, E Bravo, NT Zinner, WR Hix, K Langanke, A Mezzacappa, and K Nomoto. 2006. "Composition of the Innermost Core-Collapse Supernova Ejecta." *The Astrophysical Journal* 637(1):415-426.

Gallino R, C Arlandini, M Busso, M Lugaro, C Travaglio, O Straniero, A Chieffi, and M Limongi. 1998. "Evolution and Nucleosynthesis in Low-Mass Asymptotic Giant Branch Stars. II. Neutron Capture and the *s*-Process." *The Astrophysical Journal* 497(1):388-403.

Gamezo VN, AM Khokhlov, ES Oran, AY Chtchelkanova, and RO Rosenberg. 2003. "Thermonuclear Supernovae: Simulations of the Deflagration Stage and Their Implications." *Science* 299 (5603):77-81.

Gamezo VN, AM Khokhlov, and ES Oran. 2005. "Three-Dimensional Delayed-Detonation Model of Type Ia Supernovae." *The Astrophysical Journal* 623:337-346.

Gavai RV and S Gupta. 2003. "Pressure and Non-Linear Susceptibilities in QCD at Finite Chemical Potentials." *Physical Review D* 68:034506. DOI:10.1103/PhysRevD.68.034506. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.68.034506.

Gehrz RD, JW Truran, RE Williams, and S Starrfield. 1998. "Nucleosynthesis in Classical Novae and Its Contribution to the Interstellar Medium." *Publications of the Astronomical Society of the Pacific* 110(743):3-26.

Giannoni MJ and P Quentin. 1980. "Mass Parameters in the Adiabatic Time-Dependent Hartree-Fock Approximation. I. Theoretical Aspects; the Case of a Single Collective Variable." *Physical Review C* 21(5):2060-2075. DOI:10.1103/PhysRevC.21.2060.

Glasner SA, E Livne, and JW Truran. 1997. "Reactive Flow in Nova Outbursts." *The Astrophysical Journal* 475(2):754-762.

Glasner SA, E Livne, and JW Truran. 2007. "Novae: The Evolution from Onset of Convection to the Runaway." *The Astrophysical Journal* 665(2):1321-1333.

Göckeler M, R Horsley, D Pleiter, PEL Rakow, A Schäfer, G Schierholz, and W Schroers. 2004. (QCDSF Collaboration). "Generalized Parton Distributions from Lattice QCD." *Physical Review Letters* 92(4):042002. DOI: 10.1103/PhysRevLett.92.042002. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.92.042002.

Göckeler M, Ph Hägler, R Horsley, Y Nakamura, M Ohtani, D Pleiter, PEL Rakow, A Schäfer , G Schierholz, W Schroers, H Stuben, and JM Zanotti. 2007. (QCDSF Collaboration). "Nucleon Electromagnetic Form Factors with Wilson Fermions." *Proceedings of Science (Lattice 2007)*:038.

Göckeler M, R Horsley, T Kaltenbrunner, Y Nakamura, D Pleiter, PEL Rakow, A Schfifer, G Schierholz, H Stüben, N Warkentin, and JM Zanotti. 2008. "Nucleon Distribution Amplitudes from Lattice QCD." *Physical Review Letters* 101(11):112002.1-112002.4. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.101.112002.

Goutte H, JF Berger, P Casoli, and D Gogny. 2005. "Microscopic Approach of Fission Dynamics Applied to Fragment Kinetic Energy and Mass Distributions in 238U." *Physical Review C* 71:024316.

Gross D and F Wilczek. 1973. "Ultraviolet Behavior of Non-Abelian Gauge Theories." *Physical Review Letters* 30(26):1343-1346. DOI: 10.1103/PhysRevLett.30.1343.

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. Accessed May 10, 2009, at http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6TVN-4PJM9GJ-4&_user=2741876&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000058656&_version=1&_urlVersion=0&_userid=2741876&md5=374297aee0c21b9a20b77a7dddd07600.

Hagen G, DJ Dean, M Hjorth-Jensen, T Papenbrock, and A Schwenk. 2007b. "Benchmark Calculations for H-3, He-4, O-16 and Ca-40 with *Ab-Initio* Coupled-Cluster Theory." *Physical Review C* 76(4): Art. No. 044305. DOI:10.1103/PhysRevC.76.044305.

Hagen G, T Papenbrock, DJ Dean, and M Hjorth-Jensen. 2008. "Medium Mass Nuclei from Chiral Interactions." *Physical Review Letters* 101(9): Art. No. 092502. DOI: 10.1103/PhysRevLett.101.092502. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.101.092502.

Hägler Ph, JW Negele, DB Renner, W Schroers, T Lippert, and K Schilling. 2003. (LHP and SESAM Collaborations). "Moments of Nucleon Generalized Parton Distributions in Lattice QCD." *Physical Review Letters D* D68(3): Art. No. 034505. DOI:10.1103/PhysRevD.68.034505. Accessed May 10, at http://link.aps.org/doi/10.1103/PhysRevD.68.034505.

Hägler Ph, JW Negele, DB Renner, W Schroers, Th Lippert, and K Schilling. 2004. (LHP and SESAM Collaborations). "Transverse Structure of Nucleon Parton Distributions from Lattice QCD." *Physical Review Letters* 93(11): Art No. 112001. DOI:10.1103/PhysRevLett.93.112001. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.93.112001.

Hägler Ph, W Schroers, J Bratt, JW Negele, AV Pockinsky, RG Edwards, DG Richards, M Engelhardt, GT Fleming, B Musch, K Orginos, and DB Renner. 2008. (LHPC Collaboration). "Nucleon Generalized Parton Distributions from Full Lattice QCD." *Physical Review D* 77(9): Art. No. 094502. DOI:10.1103/PhysRevD.77.094502.

Hands S. 2007. "Simulating Dense Matter." *Progress of Theoretical Physics (Supplement)* 168:253-260. DOI:10.1143/PTPS.168.253. Accessed May 10, 2009, at http://arxiv.org/abs/hep-lat/0703017.

Hansen CJ and JC Wheeler. 1969. "A Calculation of a White Dwarf Supernova." *Astrophysics and Space Science* 3(3):464-474.

Hatsuda T. 2007. "Bulk and Spectral Observables in Lattice QCD." *Journal of Physics G: Nuclear and Particle Physics* 34:S287-S294. DOI:10.1088/0954-3899/34/8/S14. Accessed May 10, 2009, at http://stacks.iop.org/JPhysG/34/S287.

Haxton W. 2008. *Neutrino Astrophysics*. INT PUB 08-26. Accessed May 10, 2009, at http://arxiv.org/abs/0808.0735.

Herant M, W Benz, WR Hix, CL Fryer, and SA Colgate. 1994. "Inside the Supernova: A Powerful Convective Engine." *The Astrophysical Journal* 435(1, Part 1):339-361.

Herwig F, B Freytag, RM Hueckstaedt, and FX Timmes. 2006. "Hydrodynamic Simulations of He Shell Flash Convection." *The Astrophysical Journal* 642(2):1057-1074.

Hix WR, ST Parete-Koon, C Freiburghaus, and F-K Thielemann. 2007. "The QSE-Reduced Nuclear Reaction Network for Silicon Burning." *The Astrophysical Journal* 667(1):476-488.

Hoyle F. 1946. "Note on the Origin of Cosmic Rays." *Monthly Notices of the Royal Astronomical Society* 106(5):384-389.

Hoyle F and WA Fowler. 1960. "Nucleosynthesis in Supernovae." *The Astrophysical Journal* 132(3):565-590.

Hu CR and B Müller. 1997. "Classical Lattice Gauge Fields with Hard Thermal Loops." *Physics Letters* B 409:377-381.

Iancu E, A Leonidov, and LD McLerran. 2001. "The Renormalization Group Equation for the Color Glass Condensate." *Physics Letters* B 510(1-4):133. DOI:10.1016/S0370-2693(01)00524-X.

Jalilian-Marian J, A Kovner, L McLerran, and H Weigert. 1997a. "The Intrinsic Glue Distribution at Very Small x." *Physical Review D* 55:5414-5428. Accessed May 10, 2009, at http://arxiv.org/abs/hep-ph/9606337.

Jalilian-Marian, J, A Kovner, A Leonidov and H Weigert. 1997b. "The BFKL Equation from the Wilson Renormalization Group." *Nuclear Physics B* 504:415-431. Accessed May 10, 2009, at http://arxiv.org/abs/hep-ph/9701284.

Janka HT and E Müller. 1996. "Neutrino Heating, Convection, and the Mechanism of Type-II Supernova Explosions." *Astronomy & Astrophysics* 306:167.

Jansen K. 2008. "Lattice QCD: A Critical Status Report." Plenary Talk at the 26th International Symposium on Lattice Field Theory (Lattice 2008), July 14-19, Williamsburg, Virginia.

Ji X-D. 1997a. "Deeply Virtual Compton Scattering." *Physical Review D* 55(11):7114-7125. DOI:10.1103/PhysRevD.55.7114. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.55.7114.

Ji X-D. 1997b. "Gauge Invariant Decomposition of Nucleon Spin and Its Spinoff." *Physical Review Letters* 78(4):610-613.

Johnstone C and S Koscielniak. 2008. "The New Generation of FFAG Accelerators." *Division of Physics of Beams Newsletter* 12-14, American Physical Society. Accessed May 10, 2009, at http://www.aps.org/units/dpb/newsletters/upload/spring08.pdf.

Jordan GC, 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." *The Astrophysical Journal* 681(2):1448-1457.

Kajantie K, M Laine, K Rummukainen, and Y Schroder. 2003. "The Pressure of Hot QCD up to g**6 ln(1/g)." *Physical Review D* 67(10): Art. No. 105008. DOI:10.1103/PhysRevD.67.105008. Accessed May 10, at http://link.aps.org/doi/10.1103/PhysRevD.67.105008.

Karsch F and E Laemann. 1994. "Susceptibilities, the Specific Heat and a Cumulant in Two Flavor QCD." *Physical Review D* 50(11):6954-6962. DOI:10.1103/PhysRevD.50.6954. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.50.6954.

Karsch F and HW Wyld. 1985. "Complex Langevin Simulation of the SU(3) Spin Model with Nonzero Chemical Potential." *Physical Review Letters* 55:2242-2245. DOI:10.1103/PhysRevLett.55.2242. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.55.2242.

Karsch F and HW Wyld. 1987. "Thermal Green's Functions and Transport Coefficients on the Lattice." *Physical Review D* 35(8):2518-2523. DOI:10.1103/PhysRevD.35.2518. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.35.2518.

Karsch F, E Laermann, and A Peikert. 2001. "Quark Mass and Flavour Dependence of the QCD Phase Transition." *Nuclear Physics B* 605:579-599. DOI:10.1016/S0550-3213(01)00200-0.

Kercek A, W Hillebrandt, and JW Truran. 1998. "Two-Dimensional Simulations of the Thermonuclear Runaway in an Accreted Atmosphere of a C+O White Dwarf." *Astronomy & Astrophysics* 337:379-392.

Khan AA, M Göckeler, Ph Hägler, TR Hemmert, R Horsley, D Pleiter, PEL Rakow, A Schäfer, G Schierholz, T Wollenweber, and JM Zanotti. 2006. (QCDSF Collaboration). "Axial Coupling Constant of the Nucleon for Two Flavours of Dynamical Quarks in Finite and Infinite Volume." *Physical Review D* 74: Art. No. 094508. DOI:10.1103/PhysRevD.74.094508.

Khokhlov AM. 1991. "Delayed Detonation Model for Type IA Supernovae." *Astronomy & Astrophysics* 245(1):114-128.

Khokhlov AM. 1995. "Propagation of Turbulent Flames in Supernovae." *The Astrophysical Journal* 449(2, Part 1):695-713.

Kifonidis K, T Plewa, L Scheck, H-T Janka, and E Müller. 2006. "Non-Spherical Core Collapse Supernovae – II. The Late-Time Evolution of Globally Anisotropic Neutrino-Driven Explosions and Their Implications for SN 1987 A." *Astronomy & Astrophysics* 453(2):661.

Kovchegov YV. 1996. "Non-Abelian Weizsacker-Williams Field and a Two-Dimensional Effective Color Charge Density for a Very Large Nucleus." *Physical Review D* 54:5463.

Kratochvila S and P de Forcrand. 2005. "QCD at Small Baryon Number." In the 22nd International Symposium on Lattice Field Theory, June 21-26, Batavia, Illinois. *Nuclear Physics B–Proceedings Supplements* 140:514-516. DOI:10.1016/j.nuclphysbps.2004.11.148.

Lee D and T Schaefer. 2005. "Neutron Matter on the Lattice with Pionless Effective Field Theory." *Physical Review C* 72:024006.

Lee D. 2006. "Ground State Energy of Spin-1/2 Fermions in the Unitary Limit." *Physical Review B* 73(11): Art. No. 115112. DOI:10.1103/PhysRevB.73.115112. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevB.73.115112.

Levit S, JW Negele, and Z Paltiel. 1980. "Time-Dependent Mean-Field Theory and Quantized Bound States." *Physical Review C* 21(4):1603-1625. DOI:10.1103/PhysRevC.21.1603. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.21.1603.

Liebendörfer M, A Mezzacappa, F-K Thielmann, OEB Messer, WR Hix, and SW Bruenn. 2001. "Probing the Gravitational Well: No Supernova Explosion in Spherical Symmetry with General Relativistic Boltzmann Neutrino Transport." *Physical Review D* 63(10):103004-1–103004-13.

Lin DJ, A Bayliss, and RE Taam. 2006. "Low Mach Number Modeling of Type I X-Ray Burst Deflagrations." *The Astrophysical Journal* 653(1):545-557.

Lin H-W, SD Cohen, RG Edwards, and DG Richards. 2008. "First Lattice Study of the N-P_{11}(1440) Transition Form Factors." *Physical Review D* 78: Art. No. 114508. DOI:10.1103/PhysRevD.78.114508. Accessed April 28, 2009, at http://arxiv.org/abs/0803.3020.

Lisetskiy AF, BA Brown, M Horoi, and H Grawe. 2004. "New T=1 Effective Interactions for the f5/2 p3/2 p1/2 g9/2 Model Space: Implications for Valence-Mirror Symmetry and Seniority Isomers." *Physical Review C* 70: Art. No. 044314. DOI:10.1103/PhysRevC.70.044314. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.70.044314.

Litvinenko VN, I Ben-Zvi, M Blaskiewicz, Y Hao, D Kayran, E Pozdeyev, G Wang, OA Shevchenko, NA Vinokurov, GI Bell, DL Bruhwiler, AV Sobol, YS Derbenev, and S Reiche. 2008. "Progress with FEL-Based Coherent Electron Cooling." *Proceedings of the 30th FEL Conference* (in press).

Luscher M. 1986. "Volume Dependence of the Energy Spectrum in Massive Quantum Field Theories. 2. Stable Particle States." *Communications in Mathematical Physics* 104(2):177-206. DOI:10.1007/BF01211589.

MacFadyen AI and SE Woosley. 1999. "Collapsars: Gamma-Ray Bursts and Explosions in 'Failed Supernovae." *The Astrophysical Journal* 524(1):262-289.

Malitsky N, J Shah, and N Hasabnis. 2009. "EPICS-DDS." *Particle Accelerator Conference 2009*. FR5REP002. Accessed September 1, 2009, at http://www.triumf.info/hosted/PAC09/.

Mathur N, SJ Dong, KF Liu, L Mankiewicz, and NC Mukhopadhyay. 2000. "Quark Orbital Angular Momentum from Lattice QCD." *Physical Review D* 62(11): Art. No. 114504. DOI:10.1103/PhysRevD.62.114504. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.62.114504.

Matsui T and H Satz. 1986. "J/psi Suppression by Quark-Gluon Plasma Formation." *Physics Letters B* 178(4):416-422.

McLerran LD and R Venugopalan. 1994a. "Green's Functions in the Color Field of a Large Nucleus." *Physical Review D* 50:2225-2233. DOI:10.1103/PhysRevD.50.2225. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.50.2225.

McLerran LD and R Venugopalan. 1994b. "Gluon Distribution Functions for Very Large Nuclei at Small Transverse Momentum." *Physical Review D* 49(7):3352-3355. DOI:10.1103/PhysRevD.49.3352. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.49.3352.

Meakin CA and D Arnett. 2006. "Active Carbon and Oxygen Shell Burning Hydrodynamics." *The Astrophysical Journal, Letters* 637(1, Part 2):L53-56.

Meakin CA and D Arnett. 2007. "Turbulent Convection in Stellar Interiors. I. Hydrodynamic Simulation." *The Astrophysical Journal* 667(1):448-475.

Meakin CA. 2008. "Hydrodynamic Processes in Massive Stars." In *Proceedings of the International Astronomical Union*, Vol. 4, 439-449. Symposium S252, Cambridge University Press, New York, New York.

Messer OEB, SW Bruenn, JM Blondin, WR Hix, CT Lee, P Marronetti, and A Mezzacappa. 2009. "Core-collapse Supernova Simulation with CHIMERA." *Bulletin of the American Astronomical Society*, 41:466.

Mészáros P. 2006. "Gamma-Ray Bursts." *Reports on Progress in Physics* 69(8):2259-2321.

Metropolis N, A Rosenbluth, MN Rosenbluth, AH Teller, and E Teller. 1953. "Equation of State Calculations by Fast Computing Machines." *Journal of Chemistry and Physics* 21(6): Art. No. 1087. DOI:10.1063/1.1699114.

Meyer HB. 2007. "A Calculation of the Shear Viscosity in SU(3) Gluodynamics." *Physical Review D* 76(10):101701.1-101701.5. DOI:10.1103/PhysRevD.76.101701.

Meyer HB. 2008. "Energy-Momentum Tensor Correlators and Viscosity." Plenary Talk at the 26th International Symposium on Lattice Field Theory (Lattice 2008), July 14-19, Williamsburg, Virginia.

Morningstar CJ and MJ Peardon. 1997. "Efficient Glueball Simulations on Anisotropic Lattices." *Physical Review D* 56(7):4043-4061. DOI:10.1103/PhysRevD.56.4043. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.56.4043.

Morningstar CJ and MJ Peardon. 1999. "The Glueball Spectrum from an Anisotropic Lattice Study." *Physical Review D* 60(3):Art. No. 034509. DOI:10.1103/PhysRevD.60.034509.

Morningstar CJ and MJ Peardon. 2004. "Analytic Smearing of SU(3) Link Variables in Lattice QCD." *Physical Review D* 69(5): Art. No. 054501." DOI:10.1103/PhysRevD.69.054501. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.69.054501.

Mrówczynski S. 1993. "Plasma Instability at the Initial Stage of Ultrarelativistic Heavy Ion Collisions." *Physics Letters B* 314(1):118-121. DOI:10.1016/0370-2693(93)91330-P.

Müller D, D Robaschik, B Geyer, FM Dittes, and J Horejsi. 1994. "Wave Functions, Evolution Equations and Evolution Kernels from Light Ray Operators of QCD." *Fortschritte der Physik/Progress of Physics* 42(2):101-141. DOI:10.1002/prop.2190420202.

Mullowney P, P Messmer, D Smithe, and Y Choi. 2008. "Towards Kinetic Modeling of Ion Transport in an ECRIS Plasma." Presented at the Los Alamos Computer Science Symposium, Santa Fe, New Mexico.

Nagaitsev S, D Broemmelsiek, A Burov, K Carlson, C Gattuso, M Hu, T Kroc, L Prost, S Pruss, M Sutherland, CW Schmidt, A Shemyakin, V Tupikov, A Warner, G Kazakevich, and S Seletskiy. 2006. "Experimental Demonstration of Relativistic Electron Cooling," *Physical Review Letters* 96(4): Art. No. 044801. DOI:10.1103/PhysRevLett.96.044801. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.96.044801.

Nakahara Y, M Asakawa, and T Hatsuda. 1999. "Hadronic Spectral Functions in Lattice QCD." *Physical Review D* 60(9): Art. No. 091503. DOI:10.1103/PhysRevD.60.091503. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.60.091503.

Nakamura A and S Sakai. 2005. "Transport Coefficients of Gluon Plasma." *Physical Review Letters.* 94(7): Art. No. 072305. DOI:10.1103/PhysRevLett.94.072305. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.94.072305.

National Academies Press (NAP). 2003. "Connecting Quarks with the Cosmos." Board on Physics and Astronomy, National Academies Press, Washington, D.C.

National Aeronautics and Space Administration (NASA). 2007. *Laser Interferometer Space Antenna (LISA).* Accessed March 23, 2009, at http://lisa.nasa.gov.

National Research Council (NRC). 2006. *Scientific Opportunities with a Rare Isotope Facility in the United States*. National Research Council, National Academic Press, Washington, D.C.

National Science Foundation (NSF). 2008. *Advance LIGO*. Accessed May 10, 2009, at http://www.ligo.caltech.edu/advLIGO.

Navrátil P and WE Ormand. 2002. "*Ab Initio* Shell Model Calculations with Three-Body Effective Interactions for p-Shell Nuclei." *Physical Review Letters* 88(15):152502.1-152502.4.

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): Art. No. 054311. DOI:10.1103/PhysRevC.62.054311. Accessed May 10, 2009, at http://link.aps.org/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):042501.1-042501.4. DOI:10.1103/PhysRevLett.99.042501.

Nieter C and JR Cary. 2004. "VORPAL: A Versatile Plasma Simulation Code." *Journal of Computational Physics* 196(2):448-473.

Nollett KM, SC Pieper, RB Wiringa, J Carlson, and GM Hale. 2007. "Quantum Monte Carlo Calculations of Neutron-alpha Scattering." *Physical Review Letters* 99(2): Art. No. 022502.

Nomoto K, D Sugimoto, and S Neo. 1976. "Carbon Deflagration Supernova, an Alternative to Carbon Detonation." *Astrophysics and Space Science* 39(2): Art. No. L37-42.

Nomoto K, F-K Thielemann, and K Yokoi. 1984. "Accreting White Dwarf Models for Type I Supernovae. III. Carbon Deflagration Supernovae." *The Astrophysical Journal* 286(2, Part 1):644-658.

Nonaka C and S Bass. 2007. "Space-Time Evolution of Bulk QCD Matter." *Physical Review C* 75(1): Art. No. 014902. DOI: 10.1103/PhysRevC.75.014902. Assessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.75.014902.

Nordlund A. 1982. "Numerical Simulations of the Solar Granulation. I – Basic Equations and Methods." *Astronomy & Astrophysics* 107(1):1-10.

O'Hara KM, SL Hemmer, ME Gehm, SR Granade, and JE Thomas. 2002. "Observation of a Strongly Interacting Degenerate Fermi Gas of Atoms." *Science Magazine* 298(5601):2179–2182. DOI: 10.1126/science.1079107.

Okamoto M, S Aoki, R Burkhalter, S Ejiri, M Fukugita, S Hashimoto, K-I Ishikawa, N Ishizuka, Y Iwasaki, K Kanaya, T Kaneko, Y Kuramashi, V Lesk, K Nagai, M Okawa, Y Taniguchi, A Ukawa, and T Yoshié. 2002. (CP-PACS Collaboration) "Charmonium Spectrum from Quenched Anisotropic Lattice QCD." *Physical Review D* 65(9): Art. No. 094508. DOI:10.1103/PhysRevD.65.094508. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.65.094508.

Ott CD. 2008. "The Gravitational Wave Signature of Core-Collapse Supernovae." *Astrophysics* 26, 063001. Accessed September 1, 2009, at http://arxiv.org/pdf/0809.0695v2.

Ott CD. 2009. "Topical Review: The Gravitational-Wave Signature of Core-Collapse Supernovae." *Classical and Quantum Gravity* 26(6): Art. No. 063001. DOI:10.1088/0264-9381/26/6/063001. Abstract accessed May 10, 2009, at http://www.iop.org/EJ/abstract/0264-9381/26/6/063001.

Palomar Transient Factory. 2009. Accessed May 10, 2009, at http://www.astro.caltech.edu/ptf/.

Perlumutter S, G Aldering, G Goldhaber, TA Knop, P Nugent, PG Castro, S Deustua, S Fabbro, A Goobar, DE Groom, IM Hook, AG Kim, MY Kim, JC Lee, NJ Nunes, R Pain, CR Pennypacker, R Quimby, C Lidman, RS Ellis, M. Irwin, RG McMahon, P Ruiz-Lapuente, N Walton, B Schaefer, BJ Boyle, AV Filippenko, N Panagia, HJM Newberg, and WJ Couch. 1999. "Measurements of Ω and Λ from 42 High-Redshift Supernovae." *The Astrophysical Journal* 517(2):565-586.

Phillips MM. 1993. "The Absolute Magnitudes of Type Ia Supernovae." *The Astrophysical Journal, Letters* 413(2, Part 2):L105-108.

Pieper C 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.

Plewa T, AC Calder, and DQ Lamb. 2004. "Type Ia Supernova Explosion: Gravitationally Confined Detonation." *The Astrophysical Journal, Letters* 612(1, Part 1):L37-40.

Politzer H. 1973. "Reliable Perturbative Results for Strong Interactions." *Physical Review Letters* 30(26):1346-1349. DOI:10.1103/PhysRevLett.30.1346. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevLett.30.1346.

Prior C. 2007. "ICFA Beam Dynamics Newsletter 43." Accessed May 10, 2009, at http://www-bd.fnal.gov/icfabd/news.html.

Pruet J, RD Hoffman, SE Woosley, H-T Janka, and R Buras. 2006. "Nucleosynthesis in Early Supernova Winds. II. The Role of Neutrinos." *The Astrophysical Journal* 644(2):1028-1039.

Puddu G and JW Negele. 1987. "Solution of the Mean Field Equations for Spontaneous Fission." *Physical Review C* 35(3):1007-1027. DOI:10.1103/PhysRevC.35.1007. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.35.1007.

Qiang J, M Furman, and R Ryne. 2002. "Strong-Strong Beam-Beam Simulation Using a Green Function Approach." *Physical Review Special Topics−Accelerators and Beams* 5(10): Art. No. 104402. DOI: 10.1103/PhysRevSTAB.5.104402. Accessed May 10, 2009, at http://prst-ab.aps.org/abstract/PRSTAB/v5/i10/e104402.

Quaglioni S and P Navrátil. 2008. "*Ab Initio* Many-Body Calculations of n-3H, n-4He, p-3,4He, and n-10Be." *Physical Review Letters* 101(9): Art. No. 092501.

Radyushkin AV. 1997. "Nonforward Parton Distributions." *Physical Review D* 56:5524-5557. Accessed May 10, 2009, at http://arxiv.org/abs/hep-ph/9704207.

Rajagopal R and F Wilczek. 2000. "The Condensed Matter Physics of QCD." arXiv:hep-ph/0011333. Audio version available at http://online.itp.ucsb.edu/bblunch/rajagopal/.

Rapp R and J Wambach. 2000. "Chiral Symmetry Restoration and Dileptons in Relativistic Heavy-Ion Collisions." In *Advances in Nuclear Physics* 25:1, Springer, USA. DOI:10.1007/b115020.

Rauscher T and F-K Thielemann. 2000. "Astrophysical Reaction Rates from Statistical Model Calculations." *Atomic Data and Nuclear Data Tables* 75(1-2):1-351.

Ravenhall DG, CJ Pethick, and JR Wilson. 1983. "Structure of Matter Below Nuclear Saturation Density." *Physical Review Letters* 50(26):2066-2069.

RIA Working Group. 2006. "The Science of the Rare Isotope Accelerator (RIA): A Brochure from the RIA Users Community." U.S. Department of Energy, Office of Science, Washington, D.C. Accessed May 10, 2009, at http://www.orau.org/ria/history.htm.

Riess AG, AV Filippenko, P Challis, A Clocchiatti, A Diercks, PM Garnavich, RL Gilliland, CJ Hogan, S Jha, RP Kirshner, B Leibundgut, MM Phillips, D Reiss, BP Schmidt, RA Schommer, RC Smith, J Spyromilio, C Stubbs, NB Suntzeff, and J Tonry. 1998. "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant." *The Astronomical Journal* 116:1009-1038.

Romatschke P and M Strickland. 2003. "Collective Modes of an Anisotropic Quark Gluon Plasma." *Physical Review D* 68(3): Art. No. 036004. DOI:10.1103/PhysRevD.68.036004. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.68.036004.

Romatschke P and U Romatschke. 2007. "Viscosity Information from Relativistic Nuclear Collisions: How Perfect is the Fluid Observed at RHIC?" *Physical Review Letters* 99(17): Art. No. 172301.1-172301.4. DOI:10.1103/PhysRevLett.99.172301.

Röpke FK and W Hillebrandt. 2005. "Full-Star Type Ia Supernova Explosion Models." *Astronomy & Astrophysics* 431(2):635-645.

Rosswog S. 2007. "Last Moments in the Life of a Compact Binary System: Gravitational Waves, Gamma-Ray Bursts and Magnetar Formation." *Revista Mexicana de Astronomía y Astrofísica (Serie de Conferencias)* 27:57-79.

Rummukainen K and H Weigert. 2004. "Universal Features of JIMWLK and BK Evolution at Small x." *Nuclear Physics A* 739(1-2):183-226. DOI:10.1016/j.nuclphysa.2004.03.219.

Rupak G and N Shoresh. 2002. "Chiral Perturbation Theory for the Wilson Lattice Action." *Physical Review D* 66(5):054503. DOI:10.1103/PhysRevD.66.054503. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.66.054503.

Schenke B, M Strickland, A Dumitru, Y Nara, and C Greiner. 2008. "Transverse Momentum Diffusion and Jet Energy Loss in Non-Abelian Plasmas." *Physical Review C* 79(3): Art. No. 034903. DOI:10.1103/PhysRevC.79.034903. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.79.034903.

Schmidt C. 2008. "Lattice QCD at Finite Density." Plenary Talk at the 24th International Symposium on Lattice Field Theory (Lattice 2006), July 23-26, Tucson, Arizona.

Seitenzahl IR, DM Townsley, F Peng, and JW Truran. 2009. "Nuclear Statistical Equilibrium for Type Ia Supernova Simulations." *Atomic Data and Nuclear Data Tables* 95(1):96-114.

Skalski J. 2008. "Nuclear Fission with Mean-Field Instantons." *Physical Review C* 77(6): Art. No. 064610. DOI:10.1103/PhysRevC.77.064610. Accessed May 10, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.77.064610.

Smithe D. 2008. *HEL Joint Technology Office Annual Program Review*. Albuquerque, New Mexico.

Staszczak A, J Dobaczewski, and W Nazarewicz. 2005. "Skyrme–Hartree–Fock Calculations of Fission Barriers of the Heaviest and Superheavy Nuclei." *International Journal of Modern Physics E* 14(3):395-402. DOI:10.1142/S0218301305003181.

Stephanov M. 2006. "QCD Phase Diagram: An Overview." Plenary Talk at the 24th International Symposium on Lattice Field Theory (Lattice 2006), July 23-28, Tucson, Arizona.

Stetcu I, BR Barrett, and U van Kolck. 2007. "No-Core Shell Model in an Effective-Field-Theory Framework." *Physics Letters B* 653:358-362. DOI:10.1016/j.physletb.2007.07.065.

Strohmayer T and L Bildsten. 2006. "New Views of Thermonuclear Bursts." Chapter 3 in *Compact Stellar X-Ray Sources*, Cambridge Astrophysics Series 38, WHG Lewin and M van der Klis, eds., Cambridge University Press, United Kingdom.

Supernova Legacy Survey. Accessed May 10, 2009, at http://www.cfht.hawaii.edu/SNLS/.

Surman R, GC McLaughlin, M Ruffert, H-T Janka, and WR Hix. 2008. "Process Nucleosynthesis in Hot Accretion Disk Flows from Black Hole-Neutron Star Mergers." *The Astrophysical Journal, Letters* 679(2, Part 2):L117-L120.

Talon S and C Charbonnel. 2005. "Hydrodynamical Stellar Models Including Rotation, Internal Gravity Waves, and Atomic Diffusion." *Astronomy & Astrophysics* 440(3):981-994.

TechQM. 2008. *Theory Meets Experiment: Theory-Experiment Collaboration for Hot Quark Matter (TECHQM)*, RIKEN-BNL Workshop, May 6-7, Brookhaven National Laboratory, Upton, New York. Accessed May 10, 2009, at https://wiki.bnl.gov/TECHQM/index.php/Main_Page.

Timmes FX, SE Woosley, and TA Weaver. 1995. "Galactic Chemical Evolution: Hydrogen Through Zinc." *Astrophysical Journal Supplement Series* 98(2):617-658.

Tur C, A Heger, and SM Austin. 2007. "On the Sensitivity of Massive Star Nucleosynthesis and Evolution to Solar Abundances and to Uncertainties in Helium-Burning Reaction Rates." *The Astrophysical Journal* 671(1):821-827.

Umeda T, S Ejiri, S Aoki, T Hatsuda, K Kanaya, Y Maezawa, and H Ohno. 2009. "Fixed Scale Approach to Equation of State in Lattice QCD." *Physical Review D* 79(5): Art. No. 051501. DOI:10.1103/PhysRevD.79.051501.

USQCD. 2007. U.S. Lattice Quantum Chromodynamics (USQCD) Collaboration collection of white papers. Accessed September 1, 2009, at http://www.usqcd.org/collaboration.html#2007Whitepapers.

Vogel P. 2007. "Neutrino Mass and Neutrinoless Double Beta Decay." Talk at the *4th International Workshop on Neutrino Oscillations in Venice* workshop, ed. M Baldo Ceolin. Venice, Italy, April 2008. Accessed May 4, 2009, at http://arxiv.org/abs/0807.1559.

Vuorinen A. 2003. "The Pressure of QCD at Finite Temperatures and Chemical Potentials." *Physical Review D* 68(5): Art. No. 054017. DOI:10.1103/PhysRevD.68.054017. Accessed April 30, 2009, at http://link.aps.org/doi/10.1103/PhysRevD.68.054017.

Warda A, JL Egido, LM Robledo, and K Pomorski. 2002. "Self-Consistent Calculations of Fission Barriers in the Fm Region." *Physical Review C* 66(1): Art. No. 014310. DOI:10.1103/PhysRevC.66.014310. Accessed May 1, 2009, at http://link.aps.org/doi/10.1103/PhysRevC.66.014310.

Weaver TA and SE Woosley. 1993. "Nucleosynthesis in Massive Stars and the $^{12}C(\alpha, \gamma)^{16}O$ Reaction Rate." *Physics Reports* 227(1-5):65-96.

Weinberg S. 1967. "Dynamical Approach to Current Algebra." *Physical Review Letters* 18(5):188-191. DOI:10.1103/PhysRevLett.18.188.

Weinberg S. 1979. "Phenomenological Lagrangians." *Physica A* 96(1-2):327-340.

Woosley SE and TA Weaver. 1988. "Presupernova Models: Sensitivity to Convective Algorithm and Coulomb Corrections." *Physics Reports* 163(1-3):79-94.
Woosley SE, A Heger, and TA Weaver. 2002. "The Evolution and Explosion of Massive Stars." *Reviews of Modern Physics*, 74 (4): 1015-1071.

Woosley SE, S Blinnikov, and A Heger. 2007. "Pulsational Pair Instability as an Explanation for the Most Luminous Supernovae." *Nature* 450(7168):390-392.

Young G. 2008. Letter to Dr. Glenn R. Young (Director, Physics Division, Oak Ridge National Laboratory, Tennessee) from Jehanne Gillo (Acting Associate Director, DOE Office of Science for Nuclear Physics) and Michael R. Strayer (Associate Director, DOE Office of Advanced Scientific Computing), "Forefront Questions in Nuclear Physics and the Role of High-Performance Computing," September 28, 2008, Washington, D.C. Accessed September 1, 2009, at http://extremecomputing.labworks.org/nuclearphysics/references/np_charge_letter.pdf.

Zhang J, OEB Messer, AM Khokhlov, and T Plewa. 2007. "On the Evolution of Thermonuclear Flames on Large Scales." *The Astrophysical Journal* 656(1):347-365.

Zingale M, SE Woosley, CA Rendleman, MS Day, and JB Bell. 2005. "Three-dimensional Numerical Simulations of Rayleigh-Taylor Unstable Flames in Type Ia Supernovae." *The Astrophysical Journal* 632: 1021-1034.

APPENDICES

APPENDIX 1: WORKSHOP AGENDA

APPENDIX 2: WORKSHOP PARTICIPANTS

APPENDIX 3: ACRONYMS

APPENDICES

APPENDIX 1: WORKSHOP AGENDA

Monday, January 26, 2009

8:00 a.m. – 9:00 a.m.	Registration/Working Breakfast: Panel Chair Meetings		Grand Foyer
9:00 a.m. – 9:10 a.m.	Introduction: Welcome	Glenn R. Young	Salons A & B
9:10 a.m. – 9:25 a.m.	DOE-ONP Perspective	Eugene A. Henry	
9:25 a.m. – 9:40 a.m.	DOE-OASCR Perspective	Michael Strayer	
9:40 a.m. –10:10 a.m.	Major Issues in Nuclear Physics Aided by Massive Computation	David Kaplan	
10:10 a.m. – 10:35 a.m.	OASCR "Hardware into the Future"	Rick Stevens	
10:35 a.m. – 11:00 a.m.	Nuclear Forces and Cold QCD	David Richards	
11:00 a.m. – 11:15 a.m.	General Discussion		Grand Foyer
11:15 a.m. – 11:40 a.m.	Nuclear Structure and Nuclear Reactions	James P. Vary	Salons A & B
11:40 a.m. – 12:05 p.m.	Hot and Dense QCD	Frithjof Karsch	
12:05 p.m. – 12:30 p.m.	Nuclear Astrophysics	Anthony Mezzacappa	
12:30 p.m. – 12:55 p.m.	Accelerator Physics	Robert Ryne	
12:55 p.m. – 1:10 p.m.	Charge for Subpanel Sessions	David J. Dean	
1:10 p.m. – 2:00 p.m.	Working Lunch: Discussion on subpanel expectations		Grand Foyer
2:00 p.m. – 4:15 p.m.	**Subpanels**		
	Nuclear Structure and Nuclear Reactions	James P. Vary and Steven C. Pieper	Salons A & B
	Nuclear Astrophysics	Anthony Mezzacappa and George Fuller	Salon C
	Nuclear Forces & Cold QCD	Thomas Luu and David Richards	Salon D
	Hot and Dense QCD	Steffen A. Bass and Frithjof Karsch	Salon E
	Accelerator Physics	Robert Ryne	Frederick Suite
4:15 p.m. – 4:30 p.m.	General Discussion		Grand Foyer
4:30 p.m. – 5:30 p.m.	**Continue Subpanel work**		
5:30 p.m. – 6:30 p.m.	Wrap up for day (Subpanels stay together)		
6:30 p.m. – 7:30 p.m.	Working Dinner: Continue discussions as needed; compare and develop major issues and needs.		Darnestown Suite

Tuesday, January 27, 2009

8:00 a.m. – 9:00 a.m.	Working Breakfast: Summary of Day 1 and Expectations for Day 2	Glenn R. Young	Grand Foyer
9:00 a.m. – 12:30 p.m.	**Subpanels**		
	Nuclear Structure and Nuclear Reactions	James P. Vary and Steven C. Pieper	Salons A & B
	Nuclear Astrophysics	Anthony Mezzacappa and George Fuller	Salon C
	Nuclear Forces & Cold QCD	Thomas Luu and David Richards	Salon D
	Hot and Dense QCD	Steffen A. Bass and Frithjof Karsch	Salon E
	Accelerator Physics	Robert Ryne	Frederick Suite
10:15 a.m. – 10:30 a.m.	General Discussion		Grand Foyer
10:30 a.m. – 12:30 p.m.	**Continue Subpanel work**		
12:30 p.m. – 2:00 p.m.	Working Lunch: Subpanel initial discussion of recommendations		Grand Foyer
	Continue Subpanel work		
2:00 p.m. – 4:00 p.m.	Subpanel Discussion and Formulation of Recommendations		Assigned Break-Out Room
4:00 p.m. – 4:15 p.m.	General Discussion		Grand Foyer
	Continue Subpanel work		
4:15 p.m. – 5:30 p.m.	Subpanel Discussion of Recommendations		Assigned Break-Out Room
5:30 p.m. – 6:00 p.m.	Subpanel Discussion of Closeout Presentation		Assigned Break-Out Room
6:30 p.m. – 7:30 p.m.	Working Dinner: Continuation of discussions on recommendations		Darnestown Suite
7:30 p.m. – 9:00 p.m.	Discussion of recommendations – develop major recommendations for closeout		Conveyors and Organizers

Wednesday, January 28, 2009

8:00 a.m. – 9:00 a.m.	Working Breakfast: Expectations for Day 3	Glenn R. Young	Grand Foyer
9:00 a.m. – 9:25 a.m.	Subpanel Report – Nuclear Astrophysics	George Fuller	Salons A & B
9:25 a.m. – 9:50 a.m.	Subpanel Report – Accelerator Physics	Robert Ryne	
9:50 a.m. – 10:15 a.m.	Subpanel Report – Nuclear Forces and Cold QCD	Thomas Luu	
10:15 a.m. – 10:40 a.m.	Subpanel Report – Nuclear Structure and Nuclear Reactions	Steven C. Pieper	
10:40 a.m. – 11:00 a.m.	General Discussion		Grand Foyer
11:00 a.m. – 11:25 a.m.	Subpanel Report – Hot and Dense QCD	Steffen A. Bass	Salons A & B
11:25 a.m. – 11:40 a.m.	Summary Observations	Martin J. Savage	
11:40 a.m. – 12:00 p.m.	Report Schedule, Thanks	Glenn R. Young	
12:00 p.m. – 1:30 p.m.	Working Lunch: Discussion – Workshop report expectations Report Writing		Grand Foyer
1:30 p.m. – 3:30 p.m.	Discussion of Report Outline, Length, Deadlines, Boilerplate-List (Charge, Attendees, White Paper and Web Links, Agenda of Meeting, Schedule, Editorial Support)	Panel leads, organizers, conveners, and editors	Salons A & B

APPENDIX 2: WORKSHOP PARTICIPANTS

Name	E-mail Address	Affiliation
Arnett, Dave	darnett@as.arizona.edu	University of Arizona
Balantekin, Baha	baha@physics.wisc.edu	University of Wisconsin
Barnes, Ted	Ted.barnes@science.doe.gov	U.S. Department of Energy
Bass, Steffen A.	bass@phy.duke.edu	Duke University
Baxevanidis, Kyriakos	Kyriakos.Baxevanidis@ec.europa.eu	European Commission
Beane, Silas	silas@physics.unh.edu	University of New Hampshire
Beck, Micah	mbeck@cs.utk.edu	University of Tennessee
Beckman , Pete	beckman@mcs.anl.gov	Argonne National Laboratory
Bernholdt, Dave	bernholdtde@ornl.gov	Oak Ridge National Laboratory
Bertrand, Jr., Fred E.	feb@ornl.gov	Oak Ridge National Laboratory
Berz, Martin	berz@msu.edu	Michigan State University
Biven, Laura	Laura.biven@science.doe.gov	U.S. Department of Energy
Brown, Ben	Ben.Brown@science.doe.gov	U.S. Department of Energy
Brown, David	dlb@llnl.gov	Lawrence Livermore National Laboratory
Bruhwiler, David	bruhwile@txcorp.com	Tech-X Corporation
Bulgac, Aurel	bulgac@phys.washington.edu	University of Washington
Campolargo, Mario	Mario.Campolargo@ec.europa.eu	European Commission
Carlson, Joseph	carlson@lanl.gov	Los Alamos National Laboratory
Chan, Tony	tfchan@nsf.gov	National Science Foundation
Chatterjee, Lali	Lali.Chatterjee@science.doe.gov	U.S. Department of Energy
Crawford, Glen	glen.crawford@science.doe.gov	U.S. Department of Energy
De Forcrand, Philippe	forcrand@mail.cern.ch	Swiss Federal Institute of Technology Zurich
Dean, David J.	deandj@ornl.gov	Oak Ridge National Laboratory
Demmel, Jim	demmel@eecs.berkeley.edu	University of California, Berkeley
DeTar, Carleton	detar@physics.utah.edu	University of Utah
Dongarra, Jack	dongarra@eecs.utk.edu	University of Tennessee & Oak Ridge National Laboratory
Dumitru ,Adrian	Adrian.Dumitru@baruch.cuny.edu	Baruch College and RIKEN-BNL
Edwards, Robert	edwards@jlab.org	Jefferson Laboratory
Elster, Charlotte	elster@ohio.edu	Ohio University
Erdelyi, Bela	erdelyi@nicadd.niu.edu	Northern Illinois University
Fann, George	fanngi@ornl.gov	Oak Ridge National Laboratory
Foster, Ian	itf@mcs.anl.gov	University of Chicago and Argonne National Laboratory
Fryer, Chris	fryer@lanl.gov	Los Alamos National Laboratory
Fuller, George	gfuller@ucsd.edu	University of California, San Diego
Furnstahl , Dick	furnstahl.1@osu.edu	Ohio State University
Grama, Ananth	ayg@cs.purdue.edu	Purdue University
Hagen, Gaute	hageng@ornl.gov	Oak Ridge National Laboratory
Hands, Simon	s.hands@swansea.ac.uk	Swansea University
Harrod, William	William.Harrod@darpa.mil	Defense Advanced Research Projects Agency
Hatton, Angela	Angela.hatton@pnl.gov	Pacific Northwest National Laboratory
Haxton, Wick	haxton@phys.washington.edu	University of Washington

Name	E-mail Address	Affiliation
Henry, Eugene A.	Gene.Henry@science.doe.gov	U.S. Department of Energy
Hix, William	raph@ornl.gov	Oak Ridge National Laboratory
Hoang, Thuc	Thuc.Hoang@nnsa.doe.gov	National Nuclear Security Administration
Jameson, Leland	ljameson@nsf.gov	National Science Foundation
Jansen, Karl	Karl.Jansen@desy.de	DESY
Ji, Xiangdong	xji@physics.umd.edu	University of Maryland
Johnson, Calvin	cjohnson@sciences.sdsu.edu	San Diego State University
Johnson, Gary	garymichaeljohnson@gmail.com	Computational Science Solutions
Joo, Balint	bjoo@jlab.org	Jefferson Laboratory
Kaplan, David	dbkaplan@phys.washington.edu	University of Washington
Karsch, Frithjof	karsch@bnl.gov	Brookhaven National Laboratory
Kerman, Arthur	kerman@mitlns.mit.edu	Massachusetts Institute of Technology
Kewisch, Jorg	jorg@bnl.gov	Brookhaven National Laboratory
Khaleel, Moe	moe.khaleel@pnl.gov	Pacific Northwest National Laboratory
Khokhlov, Alexei	ajk@oddjob.uchicago.edu	University of Chicago
Kusnezov, Dimitri	dimitri@nnsa.doe.gov	National Nuclear Security Administration
Lamb, Don	lamb@oddjob@uchicago.edu	University of Chicago
Latham, Rob	robl@mcs.anl.gov	Argonne National Laboratory
Lee, Dean	dean_lee@ncsu.edu	North Carolina State University
Lippert, Thomas	Th.Lippert@fz-juelich.de	Julich Supercomputing Center
Liu, Keh-Fei	liu@pa.uky.edu	University of Kentucky
Livny, Miron	miron@cs.wisc.edu	University of Wisconsin
Lusk, Rusty	lusk@mcs.anl.gov	Argonne National Laboratory
Luu, Thomas	luu5@llnl.gov	Lawrence Livermore National Laboratory
Ma, Kwan-Liu	ma@cs.ucdavis.edu	University of California at Davis
MacCabe, Barney	maccabe@unm.edu	Oak Ridge National Laboratory
Marronetti, Pedro	pmarrone@fau.edu	Florida Atlantic University
Matthews, Hope	Hope.matthews@pnl.gov	Pacific Northwest National Laboratory
Meakin, Casey	Casey.meaking@gmail.com	University of Arizona
Meisner, Bob	Bob.Meisner@nnsa.doe.gov	National Nuclear Security Administration
Messer, Bronson	bronson@ornl.gov	Oak Ridge National Laboratory
Messina, Paul	messina@msc.anl.gov	Argonne National Laboratory
Meyer ,Curtis	cmeyer@cmu.edu	Carnegie Mellon University
Meyer, Harvey	meyerh@mit.edu	Massachusetts Institute of Technology
Meza, Juan	jcmeza@lbl.gov	Lawrence Berkeley National Laboratory
Mezzacappa, Anthony	mezzacappaa@ornl.gov	Oak Ridge National Laboratory
Morss, Sue	Sue.morss@science.doe.gov	U.S. Department of Energy
Mueller, Berndt	muller@phy.duke.edu	Duke University
Mustapha, Brahim	mustapha@phy.anl.gov	Argonne National Laboratory
Navrátil, Petr	navratil1@llnl.gov	Lawrence Livermore National Laboratory
Negele, John	negele@mitlns.mit.edu	Massachusetts Institute of Technology
Ng, Cho	cho@slac.stanford.edu	Stanford University
Ng, Esmond	egng@lbl.gov	Lawrence Berkeley National Laboratory
Norman, Mike	mlnorman@ucsd.edu	University of California, San Diego
Orginos, Kostas	kostas@jlab.org	Jefferson Laboratory
Ormand, William	ormand1@llnl.gov	Lawrence Livermore National Laboratory
Ostroumov, Peter	ostroumov@anl.gov	Argonne National Laboratory
Papenbrock, Thomas	tpapenbr@utk.edu	University of Tennessee

Name	E-mail Address	Affiliation
Pascucci, Valerio	pascucci@acm.org	University of Utah
Petreczky, Peter	petreczk@quark.phy.bnl.gov	Brookhaven National Laboratory
Pieper, Steven C.	spieper@anl.gov	Argonne National Laboratory
Polansky, Walt	walt.polansky@science.doe.gov	U.S. Department of Energy
Qian, Yong	qian@physics.umn.edu	University of New Mexico
Quaglioni, Sofia	quaglioni1@llnl.gov	Lawrence Livermore National Laboratory
Reddy, Sanjay	reddy@lanl.gov	Los Alamos National Laboratory
Richards, David	dgr@jlab.org	Jefferson Laboratory
Riley, Katherine	riley@mcs.anl.gov	Argonne National Laboratory
Roche, Ken	rochekj@ornl.gov	Oak Ridge National Laboratory
Rosner, Bob	r-rosner@uchicago.edu	University of Chicago
Ryne, Robert	RDRyne@lbl.gov	Lawrence Berkeley National Laboratory
Sameh, Ahmed	sameh@cs.purdue.edu	Purdue University
Savage, Martin J.	savage@phys.washington.edu	University of Washington
Schierholz, Gerrit	Gerrit.Schierholz@desy.de	DESY
Schmidt, Christian	schmidt@physik.uni-bielefeld.de	University of Bielefeld
Schroeder, Lee	lsschroeder@lbl.gov	Lawrence Berkeley National Laboratory
Schunck, Nicholas	schuncknf@ornl.gov	Oak Ridge National Laboratory
Seidel, Ed	hseidel@nsf.gov	National Science Foundation
Sekine, Yukiko	yukiko.sekine@science.doe.gov	Department of Energy
Shoshani, Arie	shoshani@lbl.gov	Lawrence Berkeley National Laboratory
Soltz, Ron	soltz@llnl.gov	Lawrence Livermore National Laboratory
Spentzouris, Pantagiotis	spentz@fnal.gov	Fermi National Accelerator Laboratory
Spotz, Bill	Bill.spotz@science.doe.gov	U.S. Department of Energy
Starrfield, Sumner	Sumner.starrfield@asu.edu	Arizona State University
Stevens, Rick	stevens@anl.gov	Argonne National Laboratory
Stoltz, Peter	pstoltz@txcorp.com	Tech-X Corporation
Straatsma, TP	tps@pnl.gov	Pacific Northwest National Laboratory
Strayer, Michael	michael.strayer@science.doe.gov	Department of Energy
Strickland, Michael	strickland@fias.uni-frankfurt.de	Gettysburg College
Teaney, Derek	derek.teaney@gmail.com	Arkansas State University
Thompson, Ian	thompson97@llnl.gov	Lawrence Livermore National Laboratory
Truran, James	truran@nova.uchicago.edu	University of Chicago
Turnbull, Bill	william.t.turnbull@noaa.gov	National Oceanic and Atmospheric Administration
Van Kolck, Ubirajara	vankolck@physics.arizona.edu	University of Arizona
Vary, James P.	jvary@iastate.edu	Iowa State University
Warren, Mike	msw@lanl.gov	Los Alamos National Laboratory
Williams, Roy	roy@cacr.caltech.edu	California Institute of Technology
Woodward, Carol	woodward6@llnl.gov	Lawrence Livermore National Laboratory
Worley, Pat	worleyph@ornl.gov	Oak Ridge National Laboratory
Young, Glenn R.	younggr@ornl.gov	Oak Ridge National Laboratory
Zacharia, Thomas	zac@ornl.gov	Oak Ridge National Laboratory
Zhang, Yuhong	yzhang@jlab.org	Jefferson Laboratory
Zhao, Quang	zhaoq@msu.edu	Michigan State University

APPENDIX 3: ACRONYMS AND ABBREVIATIONS

AG	alternating gradient
AGB	asymptotic giant branch
AMR	adaptive mesh refinement
ANL	Argonne National Laboratory
ASC	advanced simulation and computing
ASCR	U.S. Department of Energy Office of Advanced Scientific Computing Research
ATDHFB	adiabatic time-dependent Hartree-Fock-Bogoliubov (calculations)
BBN	Big Bang nucleosynthesis
BER	Office of Biological and Environmental Research
BES	U.S. Department of Energy Office of Basic Energy Sciences
BNL	Brookhaven National Laboratory
CC	coupled cluster
CEBAF	Continuous Electron Beam Accelerator Facility
CERN	European Organization for Nuclear Research
CFD	computational fluid dynamics
CGC	color-glass condensate
CI-SM	configuration-interaction shell model
CNO	carbon-nitrogen-oxygen
CP	charge-conjugation and parity
CPT	charge-conjugation, parity, and time-reversal
CPU	central processing unit
CSIRO	Commonwealth Scientific and Industrial Research Organisation
DDT	deflagration-to-detonation transition
DFT	density functional theory
DOE	U.S. Department of Energy
ECR	electron cyclotron resonance
EDM	electric dipole moment
EFT	effective field theory
EIC	electron-ion collider
ELIC	Electron Light Ion Collider
EOS	equation of state

eRHIC	Electron Relativistic Heavy Ion Collider
ERL	energy-recovery linac
EXIST	Energetic X-Ray Imaging Survey Telescope
FAIR	Facility for Antiproton and Ion Research
FEL	Free-Electron Laser (technology)
FMR	fixed mesh refinement
FRIB	Facility for Rare Isotope Beams
GCD	gravitationally confined detonation
GFMC	Green's Function Monte Carlo
GMRES-DR	Generalized Minimum Residual with Deflating Restarting (method; algorithm)
GPD	Generalized Parton Distribution
GRB	gamma-ray burst
GW	gravitational waves
HEP	high-energy physics
HFB	Hartree-Fock-Bogoliubov (calculations)
HISQ	highly improved staggered quark
HOM	high-order mode
I/O	input/output
IMF	initial mass function
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
Jefferson Laboratory	Thomas Jefferson National Accelerator Facility
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LEC	low-energy constant
LEFT	lattice-based EFT
LHC	Large Hadron Collider
LIGO	Laser Interferometer Gravitational-Wave Observatory
LLNL	Lawrence Livermore National Laboratory
LQCD	lattice quantum chromodynamics
LTE	local thermodynamic equilibrium
MPI	message-passing interface

NASA	National Aeronautics and Space Administration
NCSM	no-core shell model
NLTE	nonlocal thermodynamic equilibrium
NN	nucleon-nucleon (interaction)
NNN	three-nucleon (interaction)
NNSA	National Nuclear Security Administration
NP	U.S. Department of Energy Office of Nuclear Physics
NPLQCD	Nuclear Physics Lattice Quantum Chromodynamics (Collaboration)
NS	neutron star
NSAC	Nuclear Science Advisory Committee
NSF	National Science Foundation
NSM	neutron star merger
ODE	ordinary differential equation
ORNL	Oak Ridge National Laboratory
PES	potential energy surface
PIC	particle-in-cell
PRD	priority research direction
QCD	quantum chromodynamics
QGP	quark-gluon plasma
QMC	quantum Monte Carlo
QRPA	quasi-particle random-phase approximation
RAGE	Radiation-Adaptive Grid Eulerian (code)
RAM	random-access memory
RFQ	radio-frequency modeling
RGM	resonating group method
RHIC	Relativistic Heavy Ion Collider
RMT	Random Matrix Theory
SASI	standing accretion shock instability
SC	U.S. Department of Energy Office of Science
SciDAC	Scientific Discovery through Advanced Computing
SN	supernovae
SNe Ia	Type Ia supernovae

STAR	Solenoid Tracker at Relativistic Heavy Ion Collider
TD-SLDA	time-dependent superfluid local density approximation
TEM	transverse electromagnetic mode
UNEDF	Universal Nuclear Energy Density Functional
UTK	University of Tennessee at Knoxville
UW	University of Washington
WIMPS	weakly interacting massive particles
WP	working point
XRB	X-ray bursts

U.S. DEPARTMENT OF
ENERGY

Office of Science

Office of Advanced Scientific Computing Research
Office of Nuclear Physics

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-18739