


LARGE SCALE COMPUTING AND STORAGE REQUIREMENTS



Nuclear Physics

Report of the NERSC / NP / ASCR
Requirements Workshop
May 26 and 27, 2011

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.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, California 94720 U.S.A.



NERSC is funded by the United States Department of Energy, Office of Science, Advanced Scientific Computing Research (ASCR) program. Yukiko Sekine is the NERSC Program Manager and Ted Barnes serves as the Nuclear Physics (NP) allocation manager for NERSC.

NERSC is located at the Lawrence Berkeley National Laboratory, which is operated by the University of California for the US Department of Energy under contract DE-AC02-05CH11231.

This work was supported by the Directors of the Office of Science, Office of Nuclear Physics, and the Office of Advanced Scientific Computing Research, Facilities Division.

Large Scale Computing and Storage Requirements for Nuclear Physics Research

Workshop Report
Conducted May 26-27, 2011
Bethesda, MD

DOE Office of Science

Nuclear Physics Program Office (NP)

Office of Advanced Scientific Computing Research (ASCR)

National Energy Research Scientific Computing Center (NERSC)

Editors

Richard A. Gerber, NERSC

Harvey J. Wasserman, NERSC

Table of Contents

1	Executive Summary	1
2	DOE NP Mission	2
3	About NERSC	4
4	Workshop Background and Structure	6
5	Workshop Demographics	7
5.1	Participants	7
5.2	NERSC Projects Represented by Case Studies	8
6	Findings	10
6.1	Summary of Requirements	10
6.1.1	Scientists in NP need larger allocations of computational resources to meet key research goals in nuclear physics.	10
6.1.2	Visualization and data analytics remain a major challenge for astrophysics and accelerator simulations.	10
6.1.3	NERSC needs to continue providing and supporting the large-scale data and processing capability required for nuclear physics experiments.	11
6.2	Other Significant Observations	11
6.3	Computing Requirements	12
7	NERSC Initiatives and Plans	13
7.1	Compute Resources	13
7.2	Data	14
7.3	Visualization, Data Analytics and Databases	15
7.4	Support for Grids and Virtual Organizations	15
8	Lattice QCD for Nuclear Physics	16
8.1	Summary and Scientific Objectives	16
8.2	Methods of Solution	22
8.3	HPC Requirements	24
8.3.1	Computational and Storage Requirements Summary.....	25
8.4	Support Services and Software	26
8.5	Emerging HPC Architectures and Programming Models	26
9	Astrophysics Simulations	28
9.1	Astrophysics Simulations Overview	28
9.2	Astrophysics Simulation Case Studies	29
9.2.1	Stellar Explosions in Three Dimensions	29
9.2.2	Radiative Transfer of Astrophysical Explosions.....	34
10	Low Energy Nuclear Theory	39
10.1	Overview	39
10.2	Case Studies	39
10.2.1	Monte Carlo Calculations of Nuclei and Nucleonic Matter	39
10.2.2	<i>Ab initio</i> Nuclear Structure Calculations	44
10.2.3	Nuclear Density Functional Theory	49

11 Heavy Ion Experiments	54
11.1 Overview	54
11.2 Case Studies.....	54
11.2.1 The ALICE and STAR Heavy Ion Collider Experiments	54
12 Nuclear Physics Accelerator Research.....	61
12.1 Overview	61
12.2 Case Studies.....	61
12.2.1 Parallel Simulation of Electron Cooling Physics and Beam Transport.....	61
Appendix A. Attendee Biographies	68
Appendix B. Workshop Agenda	71
Appendix C. Abbreviations and Acronyms.....	72
Appendix D. About the Cover	74

1 Executive Summary

The National Energy Research Scientific Computing Center (NERSC) is the primary computing center for the DOE Office of Science, serving approximately 4,000 users and hosting some 550 projects that involve nearly 700 codes for a wide variety of scientific disciplines. In addition to large-scale computing resources NERSC provides critical staff support and expertise to help scientists make the most efficient use of these resources to advance the scientific mission of the Office of Science.

In May 2011, NERSC, DOE's Office of Advanced Scientific Computing Research (ASCR) and DOE's Office of Nuclear Physics (NP) held a workshop to characterize HPC requirements for NP research over the next three to five years. The effort is part of NERSC's continuing involvement in anticipating future user needs and deploying necessary resources to meet these demands.

The workshop revealed several key requirements, in addition to achieving its goal of characterizing NP computing. The key requirements include:

1. Larger allocations of computational resources at NERSC;
2. Visualization and analytics support; and
3. Support at NERSC for the unique needs of experimental nuclear physicists.

This report expands upon these key points and adds others. The results are based upon representative samples, called "case studies," of the needs of science teams within NP. The case studies were prepared by NP workshop participants and contain a summary of science goals, methods of solution, current and future computing requirements, and special software and support needs. Participants were also asked to describe their strategy for computing in the highly parallel, "multi-core" environment that is expected to dominate HPC architectures over the next few years.

The report also includes a section with NERSC responses to the workshop findings. NERSC has many initiatives already underway that address key workshop findings and all of the action items are aligned with NERSC strategic plans.

2 DOE NP Mission

The mission of the Nuclear Physics (NP) program is to discover, explore, and understand all forms of nuclear matter. The fundamental particles that compose nuclear matter - quarks and gluons - are relatively well understood, but exactly how they fit together and interact to create different types of matter in the universe is still not fully explained. To solve this mystery, NP supports experimental and theoretical research - along with the development and operation of particle accelerators and advanced technologies - to create, detect, and describe the different forms and complexities of nuclear matter that can exist in the universe, including those that are no longer found naturally.

DOE supports an extensive program of experimental research in nuclear physics at the national laboratories, which is the largest component of the Nuclear Physics effort. This experimental work is closely coupled to a strong program of research in theoretical nuclear physics, which has the goal of interpreting the experimental results in terms of our current understanding of theory, and planning future experiments and facilities to exploit advances in our understanding of the field.

In recent years, high performance computing has become remarkably important as a tool for promoting theoretical advances in nuclear physics. Although this subject originated as the study of atomic nuclei, over time the reach of nuclear physics has broadened considerably, and it now encompasses the study of all forms of strongly interacting matter. The newer topics in nuclear physics include searches for exotic phases such as the quark gluon plasma; aspects of nuclear astrophysics, including the formation of the heavier elements in supernovae, and the properties of condensed objects such as neutron stars; unusual strongly interacting meson and baryon bound states of the fundamental quarks and gluons; and the use of nuclei to search for evidence of new physics “beyond the standard model.”

A unifying theoretical theme for nuclear physics is provided by quantum chromodynamics (QCD); this quantum field theory describes the strong interaction of nuclear physics in terms of relatively simple basic interactions between the fundamental particles known as quarks and gluons. A complete solution of the equations of QCD would in principle answer most of the outstanding questions in nuclear physics. The areas of nuclear physics at the immediate frontier of applications of QCD are studies of strongly interacting particles themselves (medium energy nuclear physics) and studies of QCD phases (heavy ion collisions). Our current approach for extracting numerical predictions from this theory uses an elegant path integral method on a spacetime lattice, which relates QCD predictions to Monte Carlo integrals. As the extreme computational requirements of this approach are a limiting factor, future progress in the direct solution of QCD will depend on the scientific community having access to appropriate computing facilities.

In the original nuclear physics area of the properties of nuclei, which involves the study of many-fermion bound states of neutrons and protons, the current state of the art for extracting nuclear properties involves the use of various computational methods. These range from direct Quantum Monte Carlo studies for smaller nuclei through the traditional nuclear shell model for moderate sized nuclei to more phenomenological methods such as density functional theory for the largest nuclei. In addition to determining properties of nuclei, this area also contributes to “beyond the standard model” studies, through predictions of transition strengths that may occur under various assumptions about fundamental particle properties. A direct connection between this traditional “low energy nuclear physics” area and the fundamental theory QCD may be achieved in the relatively near future, through a realization of the “Holy Grail” of QCD-based calculations of nuclear forces.

Finally, research in nuclear astrophysics seeks to understand the nuclear processes that have shaped the cosmos — from the origin of the elements, the evolution of stars, and the detonation of supernovae, to the structure of neutron stars and the nature of matter under extreme conditions. Furthering the DOE goal of understanding the natural world, computer simulations in these areas support DOE efforts in unraveling the origins of the universe, the nature of dark energy and matter, and astrophysical production of exotic nuclei, as outlined in the 2011 DOE Strategic Plan†.

† U.S. Department of Energy Strategic Plan, May 2011

(http://energy.gov/sites/prod/files/2011_DOE_Strategic_Plan_.pdf)

3 About NERSC

The National Energy Research Scientific Computing (NERSC) Center, which is supported by the U.S. Department of Energy's Office of Advanced Scientific Computing Research (ASCR), serves more than 4,000 scientists working on over 550 projects of national importance. Operated by Lawrence Berkeley National Laboratory (LBNL), NERSC is the primary high-performance computing facility for scientists in all of the research programs supported by the Department of Energy's Office of Science. These scientists, working remotely from DOE national laboratories; universities; other federal agencies; and industry, use NERSC resources and services to further the research mission of the Office of Science (SC). While focused on research that supports DOE's missions and scientific goals, computational science conducted at NERSC spans a range of scientific disciplines, including physics, materials science, energy research, climate change, and the life sciences. This large and diverse user community runs hundreds of different application codes. Results obtained using NERSC facilities are cited in about 1,500 peer reviewed scientific papers per year. NERSC activities and scientific results are also described in the center's annual reports, newsletter articles, technical reports, and extensive online documentation. In addition to providing computational support for projects funded by the Office of Science program offices (ASCR, BER, BES, FES, HEP and NP), NERSC directly supports the Scientific Discovery through Advanced Computing (SciDAC¹) and ASCR Leadership Computing Challenge² Programs, as well as several international collaborations in which DOE is engaged. In short, NERSC supports the computational needs of the entire spectrum of DOE open science research.

The DOE Office of Science supports three major High Performance Computing Centers: NERSC and the Leadership Computing Facilities at Oak Ridge and Argonne National Laboratories. NERSC has the unique role of being solely responsible for providing HPC resources to all open scientific research areas sponsored by the Office of Science. The Leadership Computing Facilities support a more limited number of select projects, whose research areas may not span all Office of Science objectives and are not restricted to mission-relevant investigations.

This report illustrates NERSC alignment with, and responsiveness to, DOE program office needs, in this case, the needs of the Office of Nuclear Physics. The large number of projects supported by NERSC, the diversity of application codes, and its role as an incubator for scalable application codes present unique challenges to the center. As demonstrated by the overall scientific productivity by NERSC users, however, the combination of effectively managed resources and excellent user support services, the NERSC Center continues its 35-year history as a world leader in advancing computational science across a wide range of disciplines.

¹ <http://www.scidac.gov>

² http://science.energy.gov/~media/ascr/pdf/incite/docs/Allocation_process.pdf

For more information about NERSC visit the web site at <http://www.nersc.gov>.

4 Workshop Background and Structure

In support of its mission and to maintain its reputation as one of the most productive scientific computing facilities in the world, NERSC regularly collects user requirements from a variety of sources. Methods include scrutiny of the NERSC Energy Research Computing Allocations Process (ERCAP) allocation requests to DOE; workload analyses; and discussions with DOE program managers and scientist customers who use the facility.

In May 2011, the DOE Office of Advanced Scientific Computing Research (ASCR, which manages NERSC), the DOE Office of Nuclear Physics (NP), and NERSC held a workshop to gather HPC requirements for current and future science programs funded by NP. This report is the result.

This document presents a number of consensus findings. The findings are based upon a selection of case studies that serve as representative samples of NERSC research supported by NP. The case studies were chosen by the DOE Program Office Manager and NERSC personnel to provide broad coverage in both established and incipient NP research areas. Since NP supports many research endeavors in these fields the case studies presented here do not represent the entirety of NP research.

Each case study contains a description of scientific goals for today and for the future, a brief description of computational methods used, and a description of current and expected future computing needs. Since supercomputer architectures are trending toward systems with chip multiprocessors containing hundreds or thousands of cores per socket and perhaps millions of cores per system, participants were asked to describe their strategy for computing in such a highly parallel, “multi-core” environment.

Requirements presented in this document will serve as input to the NERSC planning process for systems and services, and will help ensure that NERSC continues to provide world-class resources for scientific discovery to scientists and their collaborators in support of the DOE Office of Science, Office of Fusion Energy Sciences.

NERSC and ASCR have been conducting requirements workshops for each of the six DOE Office of Sciences offices that allocate time at NERSC (ASCR, BER, BES, FES, HEP, and NP). The process began in May 2009 (with BER) and concluded in May 2011 (with NP). The target for science goals and computing requirements has been approximately 2013 for the earlier workshops and 2014 for the last two.

Specific findings from the workshop follow.

5 Workshop Demographics

5.1 Participants

Name	Institution	Area of Interest	NERSC Repo(s)
Ted Barnes	DOE NP	Nuclear Data and Nuclear Theory Computing; NP HPC Allocations	
Yukiko Sekine	DOE ASCR	NERSC Program Manager	
David Bruhwiler	Tech-X Corp	Accelerator Physics	m327
Joseph Carlson	LANL	Low Energy Nuclear Theory	m308, mp26, m1152
Robert Edwards	Jefferson Laboratory	Lattice QCD	mp7
Jon Engel	University of North Carolina	Low Energy Nuclear Theory	m308
Richard Jones	University of Connecticut	Experimental Medium Energy Nuclear Theory	
Daniel Kasen	UC Berkeley	Nuclear Astrophysics	m1186
Esmond Ng	LBNL	Low Energy Nuclear Theory, Numerical Analysis	m308, mp127
Tomasz Plewa	Florida State University	Nuclear Astrophysics	m461
R. Jeff Porter	LBNL	Heavy Ion Experiments	alice
Alejandro Sonzogni	BNL	Nuclear Data	
Martin Savage	University of Washington	Lattice QCD	m747
W.I. (Chip) Watson	Jefferson Laboratory	Lattice QCD	
Richard Gerber	NERSC	Workshop Facilitator	
Harvey Wasserman	NERSC	Workshop Facilitator	
Kathy Yelick	NERSC	NERSC Director	

5.2 NERSC Projects Represented by Case Studies

NERSC projects represented by case studies are listed in the table below, along with the number of NERSC hours they used in 2010. These projects accounted for 96 percent of time used by Nuclear Physics research at NERSC that year. (The ALICE and STAR projects described in Chapter 11 use the NERSC PDSF system, which is not allocated through the ERCAP process and is not included in this total.) The Lattice Quantum Chromodynamics (QCD) case study included in this report represents the needs of all six 2010 NP QCD projects at NERSC.

NERSC Project ID (Repo)	NERSC Project Title	Principal Investigator	Workshop Speaker(s)	Hours Used at NERSC in 2010
Lattice Quantum Chromodynamics				
m747, m789, m102, mp7, mp133, m762	<i>Hadron-Hadron Interactions with Lattice QCD</i>	Martin Savage	Robert Edwards / Martin Savage / Chip Watson	46 M
Total of projects represented by case studies				46 M
NERSC NP Lattice QCD AY2010 Total				46 M
Percent of NERSC NP Lattice QCD represented by case studies				100%
Nuclear Astrophysics				
m1186	<i>Radiation Transport of Astrophysical Explosions</i>	Daniel Kasen	Daniel Kasen	2 M
m461	<i>Supernova Explosions in Three Dimensions</i>	Tomasz Plewa	Tomasz Plewa	0.94 M
Total of projects represented by case studies				3 M
NERSC NP Astrophysics AY2010 Total				3.3 M
Percent of NERSC NP Astrophysics represented by case studies				91%
Theoretical Nuclear Physics				
m308	<i>Low-Energy Nuclear Physics National HPC Initiative: Building a Universal Nuclear Energy Density Functional</i>	Joseph Carlson	Joseph Carlson / Jonathan Engel	4.7 M
m94	<i>Structure and Reactions of Hadrons and Nuclei; Non-perturbative Solutions of Strongly Interacting Particles and Fields</i>	James Vary	Esmond Ng	3.4 M
Total of projects represented by case studies				8.1 M
NERSC Low Energy Nuclear Physics AY2010 Total				9.4 M
Percent of NERSC Low Energy Nuclear Physics represented by case studies				86%
Accelerator Physics				
m327	<i>Parallel Simulation of Electron Cooling Physics and Beam Transport</i>	David Bruhwiler	David Bruhwiler	736,616
Total of projects represented by case studies				736,616
NERSC NP Accelerators Total				819,789
Percent of NERSC Accelerators represented by case studies				90%
Heavy Ion Experiments				
ALICE	<i>Data analysis and simulations for the ALICE experiment at the LHC</i>	Peter Jacobs (LBNL)	Jeff Porter	*
m1094	<i>STAR Detector Simulations and Data Analysis</i>	Grazyna Odyniec	Jeff Porter	*
Total Hours Represented by All Case Studies				58 M
All NP at NERSC in 2010				60 M
Percent of NERSC NP Represented by Case Studies				96 %
Percent of Case Studies Represented by LQCD				80%

6 Findings

6.1 Summary of Requirements

The following is a summary of consensus requirements derived from the case studies (to be added by NERSC later). Note that many requirements are stated individually but are in fact closely related to and dependent upon others.

6.1.1 Scientists in NP need larger allocations of computational resources to meet key research goals in nuclear physics.

- a) Researchers in NP anticipate needing about five billion hours of computing time to support their research at NERSC in 2014, some 81 times more than they used in 2010.
- b) Lattice QCD simulations, which have very precise research goals and associated resource needs, require greater than 90% of the estimated NP 2014 hours.
- c) Significant NERSC resources are required to perform accelerator beam dynamics simulations that could significantly increase the luminosity of existing or future accelerators and reduce the time / cost required for obtaining important nuclear physics results.
- d) Researchers have a range of parallel processing needs that includes ensemble processing using moderate parallelisms as well as large-scale computing that requires the largest systems available.
- e) Some nuclear physics simulations, particularly low energy theoretical studies will require large memory capacity systems with a minimum of 2 GB per core in order to carry out required studies of larger nuclei.

6.1.2 Visualization and data analytics remain a major challenge for astrophysics and accelerator simulations.

- a) Visualizing four- and six-dimensional fields will be required and researchers will need help from NERSC analytics staff.
- b) Useable methods for cross-correlating across large databases that may include irregularly gridded data will be required.

6.1.3 NERSC needs to continue providing and supporting the large-scale data and processing capability required for nuclear physics experiments.

- a) Heavy ion experiment users have needs that differ from those of other NERSC users in terms of account management, scientific workflow, parallelism, disk capacity, platforms but need their processing hardware sited at NERSC to have access to other important NERSC resources, especially HPSS.
- b) Experimental scientists in NP require new NERSC user authorization procedures that rely on the PI and the Virtual Organization to manage its user base in accordance with common user agreements and practices.
- c) Some experimental scientists in NP rely on grid infrastructure, including the Open Science Grid (OSG) interface for grid submission of some production tasks.

6.2 Other Significant Observations

- a) GPU-based clusters are having a profound effect on a portion of the LQCD workload, the “analysis” phase. The increased capacity afforded by GPUs has resulted in changes in computational workflow and a significant increase in IO demands in both intermediate and long term storage.

6.3 Computing Requirements

The following table lists the AY2014 computational hours required at NERSC for research represented by the case studies in this report. “Total Scaled Requirement” at the end of the table represents the hours needed by all 2010 NP NERSC projects if increased by the same factor as that needed by the projects represented by the case studies. The 81-fold increase over 2010 NERSC use is the ratio of the sum of the hours for 2014 and 2010. Two of the Low Energy Nuclear Theory case studies, “Monte Carlo Calculations of Nuclei and Nucleonic Matter” and “Nuclear Density Functional Theory,” have their increase over 2010 NERSC use calculated relative to the same NERSC project, repository m308, “Low-Energy Nuclear Physics National HPC Initiative: Building a Universal Nuclear Energy Density Functional.”

Case Study Title	Principal Investigator	Hours Needed in 2014	Increase Over 2010 NERSC Usage
<i>Hadron-Hadron Interactions with Lattice QCD</i>	Martin Savage	4,400 M	96
<i>Stellar Explosions in Three Dimensions</i>	Tomasz Plewa	40 M	43
<i>Radiative Transfer of Astrophysical Explosions</i>	Daniel Kasen	50 M	25
<i>Monte Carlo Calculations of Nuclei and Nucleonic Matter</i>	Joseph Carlson	50 M	11
<i>Ab initio Nuclear Structure Calculations</i>	James Vary	50 M	15
<i>Nuclear Density Functional Theory</i>	Jonathan Engel	105 M	22
<i>Parallel Simulation of Electron Cooling Physics and Beam Transport</i>	David Bruhwiler	5 M	7
Total Represented by Case Studies		4,700 M	81
Percent of NERSC NP Represented by Case Studies		97 %	
All NP at NERSC Total Scaled Requirement		4,880 M	


7 NERSC Initiatives and Plans

NERSC has initiatives underway and long-term strategic plans that address some requirements presented in this report. A brief summary of these initiatives and plans is presented in this section.

7.1 Compute Resources

The NERSC Hopper system, a Cray XE6 with 1.3 PF/s of peak performance and 120 TF/s performance on a set of representative applications, was installed in the fall of 2010 and went into production on May 1, 2011. Hopper delivers a 4-fold increase in aggregate application performance over the quad-core Franklin system that went into production in mid 2009. Total allocations in 2011 and 2012 are expected to be about 1.1 billion hours (a factor of 3.5 over 2010). NERSC has started the procurement process for a NERSC-7 system, which is expected to be available for production computing by 2014 at the latest.

The figure below, showing the historical growth of NP and overall usage at NERSC, indicates that the need for computational hours to support NP in 2014 exceeds that expected by the historical trend (lower black line) and approaches the total number of hours expected at NERSC for all offices, based on current plans and funding profiles.



Total usage and NP usage at NERSC per year. The red circles represent expected number of hours at NERSC in future years, based on expected funding. The large blue X denotes the hours needed by NP at NERSC in 2014, based on this report.

There is currently great interest in using GPUs to accelerate computations. NERSC has provided users with the Dirac testbed for exploring this architectural path. While there have been successes accelerating some CPU-intensive kernels already, the NERSC user community has told us that it is not yet prepared to make a substantial investment in a technology that may not be appropriate for their entire workload. NERSC will be closely monitoring these and other low-power processor and memory technologies, communicating with vendors to help them understand the needs of the full workload, and working with the NP community to better understand ways of adapting their algorithms and software to future systems.

7.2 Data

NERSC plans to continue a constant investment in storage each year; at the planned budget levels this would result in a four-fold increase in disk capacity by 2014. Based on the findings in this report NERSC is largely on track to meet NP data requirements, although it should be noted that some QCD projects are projecting the need for 100 TB online quotas – which represents a 5-fold increase over today’s largest quotas versus our expected 4-fold projected growth.

NERSC is investing heavily in improving both capacity and bandwidth for the HPSS archival storage system. The NERSC HPSS system is designed to accommodate 50% growth per year in both total data transferred and stored. The system handled over 4 PB of IO in 2010 and grew by 2.2 PB. Additional capacity is expected to grow by a factor of two each year for the next three to five years.

NERSC is working closely with ESnet to implement the Advanced Networking Initiative-based 100Gb networks that will bring about significant improvement – 10X – in data transfer rates. NERSC supports GlobusOnline and has dedicated data transfer servers for rapid file transfers. For the foreseeable future (~five years) these capabilities are likely to represent the primary methods for data transfer between sites.

7.3 Visualization, Data Analytics and Databases

NERSC works with NP scientists to help them visualize their datasets through the NERSC Analytics / LBNL Visualization Group, whose services are available to NERSC users upon request. This group is tightly coupled with the group responsible for the VisIt interactive parallel visualization and graphical analysis tool. The collaboration helps bring new tools into the VisIt suite, assuring the ability to handle future visualization needs. VisIt is capable of visualizing complex 4 and 6-dimensional data fields, under the guidance of the Analytics Group.

Recognizing the need of the community, NERSC is in the process of inaugurating a database service that will be available to users upon request. The new service will be similar to NERSC project directories: users will request a database via a web form and will “own” their database instance in much the same way they own a project directory. There will be two new nodes that provide MySQL and Postgres instances. MongoDB is also available and is currently in use by an important materials science project.

7.4 Support for Grids and Virtual Organizations

NERSC supports the Globus software stack and Open Science Grid Compute and Storage Interfaces on all major systems, including PDSF. NERSC expects to add support for virtual organizations in 2012.

8 Lattice QCD for Nuclear Physics

Principal Investigators: Robert Edwards (Thomas Jefferson National Laboratory)
Martin Savage (University of Washington)
Chip Watson (Thomas Jefferson National Laboratory)

NERSC Repositories: m747, m789, m762, mp133, mp7

Contributors: William Detmold (College of William and Mary)
John Negele (MIT)
Frank Lee (George Washington University)
Huey-Wen Lin (University of Washington)
Keh-Feh Liu (University of Kentucky)
Kostas Orginos (College of William and Mary)

8.1 Summary and Scientific Objectives

The structure of the proton and neutron, and the forces between them, 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 us 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 been met with less success. A long-standing effort of the U.S. Department of Energy's (DOE) Nuclear Physics program is to determine how QCD in this low-energy regime manifests itself into the observed spectrum of hadrons and the observed nuclear phenomena, and to use QCD to make reliable predictions for processes that cannot be experimentally accessed.

The specific goals of Lattice QCD in nuclear physics are as follows:

1. Determine the spectrum of QCD. In addition to the excited meson and baryon spectrum, this includes the search for exotic states that may exist. This is tied closely to the experimental program at the Thomas Jefferson National Laboratory.
2. Determine how QCD makes hadrons and to quantify their structure. This is tied closely to a number of experimental nuclear physics programs, and also impacts the high-energy physics program at the LHC.
3. Determine how nuclei and their interactions emerge from QCD and develop technology to calculate these quantities with quantifiable uncertainties. This will allow for reliable predictions of nuclear reactions and structure that are difficult to

measure in the laboratory, but play a role in extreme conditions in the cosmos, or in neutron-rich environments. This includes a precise determination of the two-nucleon, three-nucleon and multi-nucleon forces.

4. Quantify the connection between the underlying fundamental symmetries of nature and experimental observables. For a number of observables that probe fundamental aspects of nature, the strong interactions provide large, and presently unquantified, modifications to the underlying interactions. Such modifications can be precisely determined from Lattice QCD calculations.
5. Understanding and quantifying the behavior of matter under extreme conditions, such as high temperature and high density. (This area is not being addressed by present work at NERSC.) The relevant sub-areas of exploration with Lattice QCD are
 - Precision calculations of bulk thermodynamics
 - QCD phase structure at non-zero net baryon number
 - Transport coefficients of QCD and spectral functions of hadrons in medium
 - Equilibration challenge: From the Color Glass Condensate to the Quark Gluon Plasma

Research into topics 1 – 4, above, is currently being conducted using NERSC resources.

Lattice QCD is the numerical technique in which space-time is discretized and the path-integral that dictates the quantum-dynamics of the quark and gluon fields is evaluated by a combination of Monte-Carlo techniques and sparse-matrix inversions. It is currently the only known method for rigorously solving QCD. Realistic calculations with uncertainties that can be systematically removed require highly optimized algorithms and cutting edge high performance computing (HPC) systems such as those that currently exist and will exist at NERSC in the near future. The current “mode of operation” in Lattice QCD is to split the requisite tasks into two or three distinct subtasks:

1. **The generation of one or more ensembles of gauge-field configurations (lattices).** Several ensembles with different lattice spacings and volumes are required at each given set of light-quark masses to perform the necessary extrapolation to the continuum and to infinite volume. This has significant resource requirements, which increase with decreasing lattice spacing, with increasing volume and with decreasing light-quark masses. Markov Chain Monte Carlo is used to produce the ensembles of independent gauge-field configurations that are saved to disk and also archived for use in (multiple) subsequent calculations. The multiple volumes and lattice spacings at a fixed quark mass allow for the reduction in the systematic error(s) introduced by the non-zero lattice spacing and the non-infinite volume of the gauge-fields. Gauge-field generation typically occurs over a period of months. The generation of gauge fields requires HPC *capability* facilities.

2. **Producing light-quark propagators on each gauge-field configuration.** A quark propagator calculated on a given gauge-field configuration encodes the probability amplitude for a quark produced at one point in space-time to be found at another point in space-time. In general, a large number of light-quark propagators need to be determined on each gauge-field configuration. There may or may not be some intermediate objects required for subsequent calculations produced at the time of propagator generation (e.g., intermediate-stage hadronic blocks). These are either saved to disk (and possibly archived) or used immediately without being written to disk. The generation of propagators presently requires HPC *capacity* facilities but can be performed equally well on *capability* resources.

3. **Correlation functions of quantities that will yield the desired physics are produced from the propagators.** They are produced immediately after propagator production or from propagators that have been saved to disk or archived. The correlation functions are written to disk, archived and moved off-site for subsequent analysis that generally can be performed on workstations or small local clusters. The generation of correlation functions require HPC *capacity* facilities but can be performed equally well on *capability* resources.

The objective of the nuclear physics program at NERSC for 2011-2014 is to perform the suite of calculations that address the science goals stated above with light-quark masses that produce a pion with the physical mass but without strong isospin breaking and without the complete electromagnetic interactions. Ensembles of gauge-field configurations will be generated with a few lattice spacings and a few lattice volumes. These, and ensembles produced elsewhere, will be used to produce the first realistic calculations of the hadron spectrum, the structure of the hadrons, and of the forces between two and three hadrons. Specifically, the requested HPC resources (through 2014) are required to perform the following calculations:

- photo-couplings of charmonium and light mesons
- cascade spectrum
- N^* spectra – initial results
- exotic meson spectrum
- isovector formfactors of the nucleon and generalized parton distributions
- individual contributions of the up, down and strange quarks to hadronic structure
- meson-meson interactions with precision
- nucleon-nucleon, hyperon-nucleon, hyperon-hyperon interactions and two-body bound-states
- light nuclei – the lowest-lying states in the lightest nuclei with $A < 8$
- neutron electric dipole moment
- constraints on physics beyond the standard model (when combined with planned experiments)

- the fine-tuning(s) of the nuclear reactions required for carbon-based life in our universe

This list of calculations can be grouped together under three specific goals (discussed above). At the physical pion mass, but in the isospin-limit and without electromagnetic interactions, the NERSC HPC resources will result in calculations of

1. the spectrum and structure of light hadrons from QCD
2. nuclear forces and the interactions between hadrons from QCD
3. the manifestation of fundamental interactions, and possible modifications to the standard model, in the light hadrons

Estimates of the computational resources that are required to complete these calculations, and calculations of other important quantities, are shown in Figure 1 and in Figure 2.




Figure 1. Estimate of the computational resources required to determine the spectrum and structure of hadrons. [Figure is reproduced from the Scientific Grand Challenges workshop Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale held in 2009]

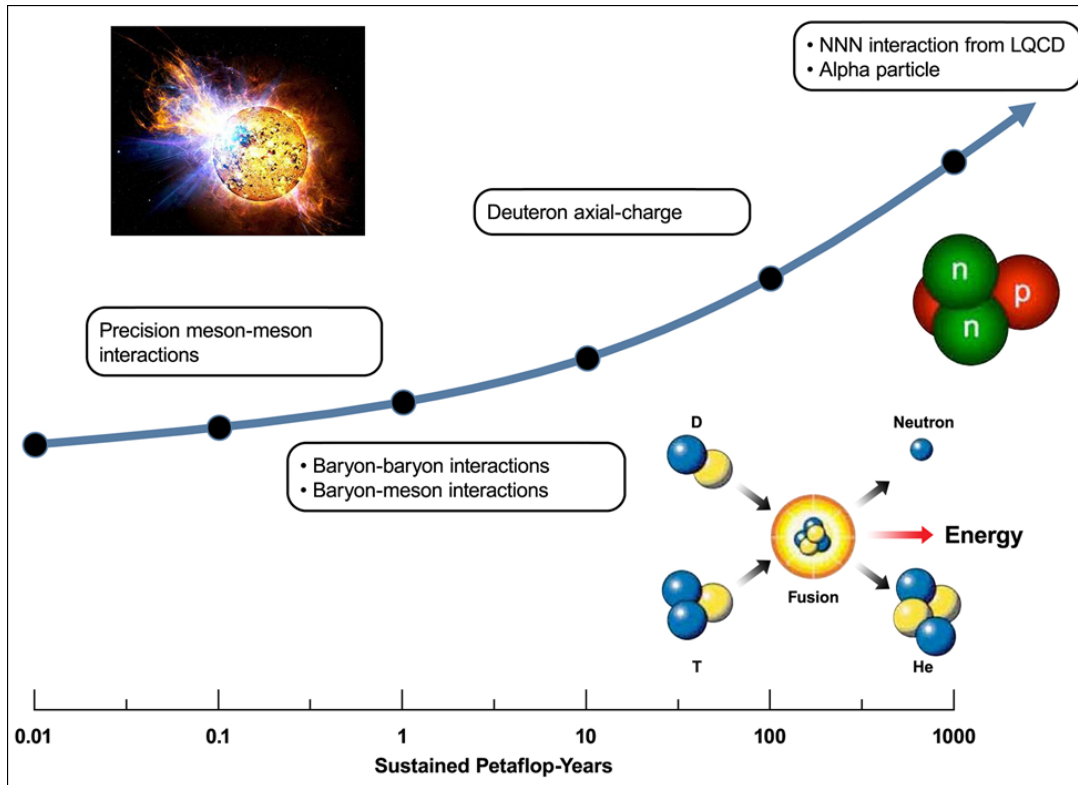


Figure 2. Estimates of the computational resources required to calculate the two and three-body nuclear forces. [Figure is reproduced from the Scientific Grand Challenges workshop Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale held in 2009]

This period (2011-2014) will also see the evolution of the currently employed Lattice QCD codes toward deployment and production on exascale facilities (workflow, etc). The resources provided by NERSC will be used coherently with other HPC resources available to the nuclear physics community, such as the *capacity* resources provided by USQCD at its HPC centers at the Thomas Jefferson laboratory and FermiLab, the *capability* resources provided by USQCD via INCITE awards, and awards from the NSF HPC centers (Teragrid).

A 50-fold increase in computational resources during the next three years over those of the previous three years, corresponding to approximately 4.4×10^9 core-hrs per year (approximately 0.5 Pflop-yrs), is required for Lattice QCD to accomplish a number of NSAC milestones during the 2011-2014 period, to complete the objectives we have outlined in section 2, and also to contribute to other key objectives of the nuclear physics program. These include the following:

The spectrum and properties of meson resonances. The presently observed spectrum of QCD provides little direct evidence of the presence of gluons. However, QCD admits 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 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 our understanding of the origin of confinement. The NSAC milestones include the 2018 milestone:

- **2018 HP15:** The first results on the search for exotic mesons using photon beams will be completed.

The structure of hadrons. LQCD will enable precision calculations of key isovector quantities, including 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. Importantly, the NSAC milestones include the 2014 milestone:

- **2014 HP9:** Perform lattice calculations in full QCD of nucleon form factors, low moments of nucleon structure functions and low moments of generalized parton distributions including flavor and spin dependence.

In addition, six of the twelve experimental milestones in *hadron physics* involve nucleon structure observables that we will address during 2011-2014:

2012 HP7: Measure the electromagnetic excitations of low-lying baryon states. (<2 GeV) and their transition form factors over the range $Q^2 = 0.1-7$ GeV² and measure the electro- and photo-production of final states with one and two pseudoscalar mesons.

Nuclei and their interactions. LQCD calculations of three- and four-baryon systems, such as the triton and alpha particle, will allow 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. For the first time, the four-nucleon system will be calculated directly from QCD. Probing this system will also provide valuable 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, currently insufficient knowledge of the three-nucleon interaction gives 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. This work will address at least two of the NSAC milestones for the period 2011-2014:

- **2014 HP10 :** Carry out *ab initio* microscopic studies of the structure and dynamics

of light nuclei based on two-nucleon and many-nucleon forces and lattice QCD calculations of hadron interaction mechanisms relevant to the origin of the nucleon-nucleon interaction.

- **Fundamental symmetries.** Preliminary calculations of the hadronic parity-violating part of the nuclear interaction will be obtained with these NERSC resources. This should coincide nicely with measurements coming from experiments at the Spallation Neutron Source at Oak Ridge National Laboratory. Relevant NSAC milestones are:
 - **2010 HI5:** Achieve factor-of-five (5) improvement in theoretical uncertainties for testing the standard model via low-energy electroweak observables.
 - **2012 HI8:** Perform independent measurements of parity-violation in few-body systems to constrain the nonleptonic weak interaction.
 - **2020 HI12:** Obtain initial results from an experiment to extend the limit on the electric dipole moment of the neutron by two orders of magnitude.

Given that the three currently distinct components of our lattice QCD calculations have quite distinct computational requirements, the increased needs will be somewhat diverse. They demand both *capability* (lattice production) and *capacity* (propagator production and correlator production). The production during the recent years has seen the distribution of CPU hours between capability and capacity in the ratio of 1:4 (approximately), and the disk requirements quite modest (most of the production output is not saved – only correlators and, in some cases, small blocks are archived). During the last year or so this distribution has been disturbed by the use of GPUs – primarily at the JLab facility. This has reduced the resource requirements for propagator production substantially (more than an order of magnitude). However, it has increased the disk requirements because more than an order of magnitude more propagators are produced that are subsequently processed on non-GPU hardware (presently). Code development is underway to port the contraction code to GPUs; however, given the nature of the contractions that must be performed, the performance that may be achieved for this part of the calculation with GPUs is unknown.

8.2 Methods of Solution

In 2010 the “Chroma” lattice field theory code suite was ranked as the 3rd code in terms of allocated cycles at NERSC and will be discussed here. This code suite, used for lattice QCD calculations, runs on almost all parallel machines. It was designed from the ground up as part of the DOE SciDAC effort within the USQCD initiative. The Chroma software system is built over a C++ data parallel API and implementation called QDP++. This package provides an architectural independent programming API along with IO support. The Chroma suite is designed following modern software engineering programming practices with regression tests and nightly builds. Significant effort has gone into optimizing architecturally specific, time critical routines.

Lattice QCD uses a four dimensional rectangular grid to represent space-time with the grid spacing and size being input parameters of the calculation. The grid is divided into equal domains for each processing node. The quarks are site variables, while the “gluons” – the field configuration – are situated on the links. For a given set of input parameters of the theory; namely the lattice spacing, volume and quark masses, an observable is computed by averaging over all possible values of the field configurations. A Monte Carlo method is used to produce the most important set of configurations contributing to this average. These importance sampling methods allow for a systematically improvable calculation of an observable.

In broad terms, there are two main classes of computations – the generation of the gauge field configurations requiring large-scale capability resources, and the analysis of these configurations requiring either *capability* or *capacity* resources. The “Chroma” code suite is actually a library with parts of the code supporting both of these types of calculations.

Field configurations are generated by a molecular dynamics evolution of the gauge field through an artificial simulation time. The time consuming part of the calculation is the computation of the force coming from the dynamical quark fields, which is a non-local force involving the gluon fields. Computations of this force require the solution of large sparse linear systems of equations; namely, the solution of a Hermitian positive-definite sparse matrix problem.

Significant effort, as a part of the USQCD SciDAC initiative, has gone into better mathematical formulations of the gauge updating – a Hamiltonian integrator. The formulation of QCD, and the demands that the integrator satisfy reversibility and be area-preserving, limit the extent to which improved mathematical formulations can be borrowed from standard ODEs. Nevertheless, over the last few years, improved methods have resulted in about a factor of ten increase in the performance of the integrator.

Once the gauge field configurations are computed and stored, subsequent calculations of observables will typically be dominated in cost by repeated solution of these matrix problems. Some calculations may involve the intermediate storage of large files scaling with the lattice volume. These files are later staged back in and processed into numerous, but smaller files.

Lattice grid sizes in use now range up to $40 \times 40 \times 40 \times 256$. The gauge field variables on each link are 3×3 complex matrices, and the site variables – the quarks - are represented as 3×4 complex matrices.

The Chroma codes can be compiled in several architectural modes supporting a hybrid communications and threads model. The parallel version is built over a communications package called QMP that was developed as a part of the USQCD SciDAC effort. This communication package is implemented in MPI as well as other architecturally specific hardware communication variants. The thread model is implemented in either OpenMP

as well as a package called QMT that was also developed as a part of SciDAC. The QMT package has been implemented in Pthreads with architecturally specific thread semaphore routines that are found to outperform OpenMP on some platforms. The choice of computational model is usually suggested by the hardware architecture.

Gauge generation on the largest grid size, $40^3 \times 256$, utilizes 40k cores on the JaguarPF system at ORNL. The hybrid QMP/QMT model, implemented as MPI/threads, is essential to achieve the best performance. Further analysis jobs are using 10k and 20k cores on the XE6 at NERSC. Gauge generation on the Bluegene/P at ANL typically uses 32k cores.

The deployment of large GPU-based commodity clusters at many sites, including the large cluster at Jefferson Lab within the USQCD facilities project, is having a profound impact on the analysis phase of lattice QCD calculations. The Jefferson Lab GPU cluster has about 100 Teraflops of sustained performance on lattice QCD benchmark. This resource, mostly utilized for matrix problem solutions, equals that available to all of lattice QCD within the US on leadership class machines. Many of the jobs are using a parallel version of a matrix solver, running typically over 4 to 16 GPUs, with 4 GPUs per compute node. This shift to using GPUs has impact on lattice QCD workflows, and groups have to restructure accordingly. As DOE leadership facilities deploy GPU clusters, it is expected that their usage as a *capacity* based resource will increase.

The increased *capacity* afforded by GPUs has resulted in a significant increase in IO demands in both intermediate and long term storage. Intermediate datasets are approaching 1 TB. Parallel and global file systems available to hold longer term datasets, are being used. A common such system is Lustre, but other HPSS systems can be used.

8.3 HPC Requirements

Lattice QCD calculations at the physical pion mass limit are required to achieve the scientific milestones of nuclear physics outlined above. They will remove the dominant systematic error resulting from the extrapolation in the light quark masses that is present in current calculations. The 4.4 B core-hour resource estimate shown in Figure 3 as the lower red (darker) star, corresponds to the production of an ensemble of 1,000 gauge-fields with physical volume of $6^3 \times 12 \text{ fm}^4$ and its subsequent analysis. The cost for generating this gauge ensemble is 0.1 PF-yr, and with the subsequent analysis at a 1:4 relative resource requirement, the total requirement will be ~ 0.5 PF-yr, or approximately 4.4 B core-hours.




Figure 3: Estimates of compute resources requirements to generate gauge field configurations circa 2009. The lower red (darker) star shows the resources required to generate a single ensemble of gauge-field configurations and perform the necessary analysis. The upper red (lighter) star denotes the cost of generating and analyzing a second ensemble with a smaller lattice spacing.

Removing the systematic error resulting from the finite lattice spacing will require calculations at more than one lattice spacing, along with an extrapolation to the continuum. The generation of one ensemble at the selected second lattice spacing, and its subsequent analysis, will require approximately 44 B core-hours, as shown by the upper (lighter) red star in Figure 3. The continuum extrapolation will parametrically reduce the lattice-spacing error, and can be further refined by calculations in further lattice ensembles. The effects of the finite lattice volume are exponentially suppressed with lattice size and are predicted to be small in these ensembles.

8.3.1 Computational and Storage Requirements Summary

	Current (NERSC 2010)	2014 Requirement
Computational Hours	73 M core-hours	4.4 B core-hours
Parallel Concurrency in typical production run	>50 K	~ 128 K
Wall Hours per Run	12-24	12-24
Aggregate Memory	1 TB	2 TB
Memory per Core	< 1.0 GB	< 1.0 GB
I/O per Run Needed	10 GB (input) and 200 GB (output)	50 GB (input) and 5 TB (output)
On-Line Storage Needed	20 TB and 100K files	100 TB and 1K files
Archival Storage Needed	100 TB	0.5 PB and 1M files
Data Transfer Needed	1 TB per month	2 TB per month

8.4 Support Services and Software

USQCD codes rely on the LAPACK and BLAS software libraries, but not much beyond these. Useful services can extend to training workshops or NDA access to new emerging architectural systems. Typically, there is not a need for significant amounts of consulting or training.

On-line disk storage and long-term tape storage demands are increasing with more ambitious analysis campaigns. Access to global file systems is a benefit, as well as the short-term storage of large datasets on the order of a few TB.

8.5 Emerging HPC Architectures and Programming Models

The development of a domain-decomposed formulation of the gauge field integrators as well as inverters is precisely aimed at large scale heterogeneous systems. Crucial to these methods is the importance of the characteristic length scale within QCD. This length scale, basically the distance for confinement of quarks within QCD, indicates the minimal size by which domains of space-time can be decomposed and isolated from other domains. These domains of space can be allocated to a computational node with fast local communications (say through memory), but relatively slow off-node communications. This situation is quite characteristic of GPU embedded accelerators.

The demonstration of the high performance of a simple Schwarz-based decomposed solver on a commodity GPU cluster scaling up to 256 GPUs with a total sustained performance of ~ 20 Teraflops (equivalent to $\sim 25K$ cores on Hopper) is a strong indication that such decomposition techniques are appropriate. More refined decomposition methods, involving multi-grid based methods, can be expected to yield even more improvements in performance. These developments will significantly improve analysis job performance and throughput.




Figure 3. *Left: sustained strong scaling performance in GFLOP/s of a Wilson-clover mixed precision BiCGstab and GCR-DD linear system solvers. Right: sustained strong scaling time to solution in seconds for the same solvers. [Figure is reproduced Scaling Lattice QCD beyond 100 GPU-s, authors R. Babich, M. Clark, B. Joo, G. Shi, R. Brower, S. Gottlieb, submitted to SC11.]*

To best utilize such large scale heterogeneous systems for gauge generation, further mathematical and algorithm developments with the concomitant software development are required. The Schwarz-based decomposed algorithm has been extended to gauge generation methods as shown by Luscher. These techniques will be incorporated into the USQCD gauge generation code suites; namely, into Chroma, and the hope is that gauge generation throughput will be substantially improved. The GPU based commodity clusters that are available today, provide a test-bed for such development efforts. In addition, the development of a heterogeneous aware data-parallel computing interface (QDP++), is specifically targeted at such large-scale systems, and can be used to further exploit the capabilities of such architectures.

9 Astrophysics Simulations

9.1 Astrophysics Simulations Overview

Nuclear physicists seek to understand how protons and neutrons combine to form atomic nuclei and how these nuclei have arisen during the 13.7 billion years since the birth of the cosmos. The forces that bind protons and neutrons together into nuclei are immensely strong, with the result that nuclear processes such as nuclear fusion and fission can release huge amounts of energy. Nuclear scientists working in astrophysics seek to understand the processes that have shaped the cosmos, from the origin of the elements, the evolution of stars, and the detonation of supernovae, to the structure of neutron stars and the nature of matter at extreme densities. Nuclear scientists have taken great strides in nuclear astrophysics, for example by decreasing the limits of the age of the universe by about one billion years through studies of the reaction cross sections that control hydrogen burning in stars. The past decade has produced major advances in our knowledge of the evolution of the universe and the origin of the elements. With these advances, we now know that nuclear reactions on unstable nuclei are extremely important in astrophysical environments. In some explosive situations, such as a supernova, they dominate. Key questions of interest include:

- What is the nature of the nuclear force that binds protons and neutrons into stable nuclei and rare isotopes?
- What is the origin of simple patterns in complex nuclei?
- What is the nature of neutron stars and dense nuclear matter?
- What is the origin of the elements in the cosmos?
- What are the nuclear reactions that drive stars and stellar explosions?

9.2 Astrophysics Simulation Case Studies

9.2.1 Stellar Explosions in Three Dimensions

Principal Investigator: Tomasz Plewa (Florida State University)
NERSC Repository: m461

9.2.1.1 Summary and Scientific Objectives

The theory of stellar evolution of intermediate mass and massive stars predicts that such objects typically end their lives as supernovae: some of the brightest objects observed in the universe. The most massive stars tend to explode following a collapse of their central regions, while others produce supernovae in the process of thermonuclear incineration of their strongly degenerate cores. These stellar explosions involve highly nonlinear phenomena. Studying them ultimately requires development of computer models and execution of large-scale computer simulations. This project is aimed at obtaining supernova explosion model observables for both core-collapse and thermonuclear events to test predictions of stellar evolution theory and gain insights into extremely complex basic physics processes.

In contrast to thermonuclear supernova modeling, no more than perhaps one dozen three-dimensional core-collapse models exist. The primary difficulty here is the intrinsic complexity of the supernova physics. Computational time used by an explosion model is typically dominated by modeling the transport of neutrinos emitted by the forming neutron star and their interactions with the infalling stellar envelope. A host of relevant hydrodynamic instabilities are known at present, including the standing accretion shock instability (SASI), and neutrino-driven convection, both of which occur during about the first second of a supernova explosion, and which are followed by a subsequent, second round of instabilities of the Rayleigh-Taylor, Richtmyer-Meshkov and Kelvin-Helmholtz types on somewhat longer time scales. The best-studied example in this respect is supernova SN 1987A.

Multidimensional models that include both a neutrino-driven initiation of the explosion and the following shock-envelope interaction, and are evolved into the phase of homologity and beyond, when they can be more directly compared to observations, are lacking completely. A major goal of this project is to provide this type of modeling in three spatial dimensions in order to account for both the most relevant physics and the proper geometry of the flow - from the early phases after core bounce until the early SN remnant phase. In particular, in regard to SN 1987A, no successful model matching both explosion energetics and spectral observations (high-velocity of radioactive material) exists. We expect our current series of models to bring theory and observations into closer agreement to help identify physics process(es) essential for explaining observations.

9.2.1.2 Methods of Solution

The supernova explosion models are obtained with the FLASH [1] code with the Proteus library of physics solvers developed for efficient modeling in the area stellar astrophysics applications. The FLASH code was developed and has been maintained at the University of Chicago since the late 1990s. FLASH is fundamentally a 3-D, Eulerian code with an adaptive block-structured octree grid. The mesh is based on the PARAMESH package. The primary hydro solver of FLASH is based on the HOTB/PROMETHEUS supernova code and uses the Piecewise Parabolic Method (PPM) to solve the compressible, inviscid fluid flow equations. PPM is a high-resolution shock-capturing Godunov-type scheme. FLASH provides support for self-gravity and advection of multiple nuclear species. It also supports several nuclear reaction networks and includes realistic, stellar equation of state. FLASH has been successfully applied to several astrophysical problems [2-5]. Preliminary validation of the code has been done in the context of the Rayleigh-Taylor (RT) instability [5-6] and shock-induced vorticity-dominated flows [7]. Our recent Proteus/FLASH applications include 3-dimensional simulation studies of a single- and multi-mode Rayleigh-Taylor experiment (the Omega laser, Rochester) and 2-dimensional simulation studies of the Kelvin-Helmholtz instability (the Nike laser, NRL).

The Proteus library offers improved resolution of the hydrodynamics module [50], an advanced multi-species interpolation scheme, and an improved meshing package for efficient support of multigrid solvers (self-gravity, thermal transport). Of particular importance to this project is the adaptation of the Hybrid Characteristics (HC) algorithm [8], originally designed for and applied to photon transport, for neutrino transport. The combination of Proteus and HC will allow us to model 3D situations in Cartesian coordinates, something that is not possible with our original supernova code. This approach offers several advantages over simulations performed on spherical meshes. We can directly control mesh resolution in space and time (thanks to the adaptive mesh refinement), avoid severe time step restrictions (due to the lack of an angular coordinate), avoid numerical artifacts that pollute the solution near the symmetry axis when spherical coordinates are used (see, e.g., [9-11]), and the neutron star (central point mass) can freely move on the grid. Thanks to the scalability of the FLASH particle module, we are able to seed simulations with a large number of tracer particles to obtain detailed distributions of nucleosynthetic products and yields.

- [1] Fryxell, B., Olson, K., Ricker, P., et al. 2000, *ApJS*, 131, 273
- [2] Alexakis, A., Calder, A. C., Heger, A., et al., 2004, *ApJ*, 602, 931
- [3] Plewa, T., Calder, A. C., & Lamb, D. Q. 2004, *ApJ*, 612, L37
- [4] Plewa, T. 2007, *ApJ*, 657, 942
- [5] Calder, A.C., Fryxell, B., Plewa, T., et al. 2002, *ApJS*, 143, 201
- [6] Dimonte, G., Youngs, D. L., Dimits, A., et al. 2004, *Phys. Fluids*, 16, 1668
- [7] Weirs, V. G., Dupont, T., & Plewa, T. 2008, *Phys. Fluids*, 20, 044102
- [8] Rijkhorst, E.-J., Plewa, T., Dubey, A. & Mellema, G. 2006, *A&A*, 452, 907
- [9] Scheck, L., Kifonidis, K., Janka, H.-T., & Müller, E. 2006, *A&A*, 457, 963

- [10] Hammer, N. J., Janka, H. ., & Mueller, E. 2009, ArXiv e-prints
[11] Scheck, L. 2007, PhD thesis, Technische Universitaet Muenchen

9.2.1.3 HPC Requirements

Modeling of the supernova explosion phase, the post-explosion envelope mixing, and the subsequent evolution toward homology is an extremely demanding computational task. The cost is primarily determined by the hydro mesh resolution required to accurately capture the dynamics of the neutrino driven convection and various fluid mixing instabilities present in both the post-explosion phase and the evolution toward a supernova remnant.

The current series of models typically use about 1,000 compute cores per model. Computations are CPU bound and the number of cores is limited primarily by the size of our computer time allocation. The computational time is dominated by hydrodynamics, which accounts for about 75% of the total. The above characteristics apply to all phases of the supernova evolution. A single, three-dimensional model of the shock-envelope interaction requires about 20,000 time steps. At a cost of about 30 seconds per time step on 1,000 cores, the total cost of obtaining such a model is approximately 150,000 compute hours. Additional nucleosynthesis calculations are embarrassingly parallel, and use several thousand cores with a typical wall-clock time of a few hours per run.

The FLASH code's I/O routines use a parallel or, if necessary, a serial version of the HDF5 library. The restart file (FLASH checkpoint) contains all the necessary information required for restarting. All processors participate in writing a restart file. HDF5 collective I/O offers modest improvements (~ 20%). The storage requirements for the explosion and post-explosion phases are similar and include several tens of restart files, each approximately 50 GB in size, and several hundred output files with size reduced to about 10 GB each. The total required storage for one phase of the supernova evolution is therefore approximately 10 TB per one complete model. Most storage consists of files with sizes exceeding 1 GB. Files are typically accessed in one of two modes: (1) in complete sequences with short lifetimes outside the archive (e.g. for visualization), and (2) focused analysis around certain evolutionary times with long lifetimes outside the archive. We extensively use NERSC Global Filesystem project data area for data analysis and visualization partitions.

Doubling the mesh resolution in the 2014 series of computations is a reasonable goal. The increased resolution will allow for better resolution of mixing instabilities, improved coupling between various physics processes, and improve accuracy of nucleosynthesis calculations. After factoring in savings due to mesh adaptivity and assuming similar performance per core, models with the increased resolution will require at least ten times more cores (10,000; proportionally more if the amount of memory available per core is reduced), ten times more computing time (1,000,000 hours), and ten times larger storage (100 TB per model). Computations will remain CPU-bound with memory bandwidth, the

speed of interconnect, and load imbalance inherent for certain physics solvers (e.g., multigrid solver for Poisson equation) being primary performance bottlenecks.

9.2.1.4 Computational and Storage Requirements Summary

The numbers in the following table assume a computational campaign producing a set of 20 supernova models.

	Current (NERSC 2010)	2014 Requirement
Computational Hours	2 M	40 M
Parallel Concurrency in typical production run	1 k	30 k
Wall Hours per Run	100,000	2,000,000
Aggregate Memory	150 GB	2 TB
Memory per Core	500 MB	150 MB
I/O per Run Needed	5 TB	50 TB
On-Line Storage Needed	500 GB	5 TB
Archival Storage Needed	100 TB	1 PB
Data Transfer Needed	500 GB	5 TB

9.2.1.5 Support Services and Software

Performance characteristics of the FLASH code are known to depend on the (Fortran) compiler and vary by about 20%. This is substantial given that a typical runtime per model is several days. It is therefore important to have access to a broad spectrum of compilers. This also aids in code development, including implementation verification and dynamic testing, as the compiler diagnostic and the range of supported language features (especially for relatively fast-evolving Fortran) vary from compiler to compiler. Availability of GPGPU-capable compilers will be very important in porting applications to the future, memory-limited computing platforms. Memory diagnostics tools such as Valgrind proved indispensable in eliminating use of uninitialized variables and memory leaks.

The FLASH code is instrumented with MPI-based timers that offer good insight into overall performance of the application, especially its parallel characteristics. However, it is expected that future HPC systems will offer not only significant less memory available to the individual process, but the memory access characteristics will likely substantially differ from the current systems. This will make a single-process analysis more important, especially in regard to memory access and memory bandwidth use. Availability of tools helping in this analysis will be very important.

The ultimate product of simulations is a synthetic image of the physical system evolution. This requires efficient visualization and data analysis tools, fast and reliable access to data itself, including data storage. Due to the extreme size of simulation data, tools must operate remotely. Select small subsets of data (~1%) will require more careful analysis and must remain available online for extended periods of time, which means accessible until at least the next generation of models is available. This is important for both model

development and assessment of the results, a cycle that typically lasts about two to three years. Again, this is for only about one percent of the data, an amount that is sufficient to spot any qualitative differences and gain some insight into quantitative changes. Additionally, these files must be accessible to all collaborators in the project (via shared file systems such as NERSC's /project file system). On the other hand, data visualization typically requires instantaneous access to all data from a single model for a limited amount of time (similar to scratch). This suggests the existence of two different file systems accessible across various computing platforms: (1) persistent and shared between subgroups of users (~10 TB in 2014), and (2) temporary and with access limited to individual team members (~50 TB in 2014).

9.2.1.6 Emerging HPC Architectures and Programming Models

The physics solvers such as ray-tracing module and nuclear networks contain substantial amount of computations that can be done in isolation of other nodes and thus offer good potential for using shared memory OpenMP-like programming models. Although OpenMP implementations typically require less programming effort, it remains unclear whether they offer substantial gains over MPI-based codes when it comes to application performance. Such information could be obtained by simulating application behavior on anticipated future systems (i.e. reduced per-process memory) along with the single process memory use characteristics (i.e. access to the main memory, cache reuse, data locality, etc.). Expert assistance in such analysis would be very helpful to the code and application developers.

9.2.2 Radiative Transfer of Astrophysical Explosions

Principal Investigator: Daniel Kasen (UC Berkeley / LBNL)

Contributors: Ann Almgren, Weiquin Zhang, and John Bell (LBNL)
Stan Woosley (UC Santa Cruz)

NERSC Repositories: m1186, m106

9.2.2.1 Summary and Scientific Objectives

The heavier elements in the universe were created and dispersed in the supernova explosions of stars. Simulating these events involves a fascinating interplay of hydrodynamics, nuclear reactions, neutrino physics, photon transport, and radioactive decay. Astronomical observations provide a rich and constraining data set, but our theoretical understanding of supernovae (SNe) remains incomplete. In particular, the progenitor stars of thermonuclear (Type Ia) SNe are still unknown, while the robustness of the neutrino mechanism for core collapse supernovae (CC SNe) has not yet been demonstrated. The astrophysical origin of the heavy R-process nuclei remains uncertain, but may be related to the energetic merger of two neutron stars and/or CC SNe. Given the complexity of these phenomena, numerical modeling in three spatial dimensions (3-D) is required to address the fundamental questions. Our scientific objective for the next few years is to advance high-fidelity simulations of astrophysical explosions, and use them to connect the nuclear theory of SNe to current and future observational and experimental programs.

This case study focuses on one aspect of supernova modeling in particular – the calculation of the transport of radiation, both electromagnetic and neutrino. (See the accompanying case study by Tomek Plewa for related issues in adaptive mesh refinement (AMR) hydrodynamics and reactive flows). The radiation transport problem is fundamental to supernova dynamics and observables in several ways: (1) In CC SNe, the energy transported by neutrinos likely drives the explosion itself, while also producing an observable neutrino burst that probes unknown neutrino physics; (2) The absorption of neutrinos around the core profoundly influences the final nucleosynthetic yields; (3) After the explosion, gamma-rays from radioactive decay propagate through and heat the remnant; (4) Finally, the radioactively heated remnant emits optical radiation, which diffuses out and produces the supernova light curves and spectra that we routinely observe on earth, and which can be used to test our predictions for the creation of the elements within these events.

We will discuss two codes implementing radiation transport methods. The CASTRO code (Almgren et al., 2010, Zhang et al., 2011) is designed to model the dynamics of astrophysical explosions. It couples compressible AMR hydrodynamics to nuclear

burning and a radiation transport module based on multi-group flux limited diffusion (MGFLD). The SEDONA code (Kasen et al., 2007) uses implicit Monte Carlo methods to post-process explosion simulations and calculates synthetic light curves and spectra that can be compared to observations. Our goal is to develop an end-to-end code framework that can follow the supernova event from the first moments of ignition/collapse, to the dynamics of the explosion, and through the subsequent evolution of the expanding, radiating remnant. This task requires the seamless integration of our different codes treating the different phases of the event.

The possible applications of the CASTRO and SEDONA codes are numerous; to focus the discussion here, we consider two specific objectives for the coming three years:

(1) Use the CASTRO code to carry out high-resolution 3-D simulations of CC SNe using MGFLD to treat the neutrino transport. This will address the robustness of the neutrino driven explosion mechanism (which is still in question) and will allow us to make predictions as to the energetics and nucleosynthetic yields of CC SNe.

(2) Use the SEDONA code to calculate the (photon and neutrino) light curves and spectra for a range of 3-D supernova models (produced by CASTRO or similar codes). We plan to consider core-collapse supernovae, as well as thermonuclear supernovae and neutrino star mergers. Comparing the light curve calculations to observations will provide the ultimate means of validating/falsifying the simulations and testing their predicted nucleosynthesis, neutrino physics, and the underlying paradigms for the explosion mechanism and progenitor stars.

9.2.2.2 Methods of Solution

SEDONA is a multi-dimensional time-dependent multi-wavelength radiation transport code that calculates the light curves and spectra of supernovae and other explosive transients of all types. The code uses implicit Monte Carlo methods to model the propagation and thermalization of gamma rays from radioactive decay and the subsequent reemission and transport of optical photons. The radiation transport is coupled to the thermodynamics of a homologously expanding supernova remnant. The code uses extensive atomic data as input to calculate detailed opacities from gamma rays to the infrared. Calculations can be performed in one, two or three spatial dimensions using various coordinate systems (Cartesian, cylindrical, spherical) as well as using adaptive grids.

SEDONA is coded in C++ and uses hybrid MPI/OpenMP parallelization. For problems with modest memory requirements (2-D or low resolution 3-D simulations) the computational domain can be fully replicated in each MPI process, which restricts communication to a few global quantities that are reduced at deterministic times. This approach has allowed for nearly perfect weak and strong scaling out to more than 100,000 cores. For high-resolution calculations, domain decomposition is required, which necessitates communication of radiation packets when they reach local grid

boundaries. Recently, we have explored domain decomposed Monte Carlo transport within the BoxLib AMR framework and found (for a uniform density test problem) excellent weak scaling out to 50,000 cores on Hopper. We found the best scaling results using OpenMP threading on each 6-core NUMA node and MPI communication between NUMA nodes. We also use spatial domain decomposition to parallelize the calculation of the matter thermodynamic state and atomic opacities, which in most runs comprises nearly half of the total computational load.

CASTRO is a compressible Eulerian radiation hydrodynamics code that is used (among other things) for simulating the dynamics of supernova explosions. The code includes self-gravity and reaction networks, and supports radiation transfer in the multi-group flux-limited diffusion (MGFLD) approximation. The adaptive refinement in CASTRO uses the BoxLib package developed by LBNL by CCSE. The parallelization strategy for the hydrodynamics is to distribute grids to processors using MPI and to use OpenMP threading to solve the hydrodynamics on the local domain. An implicit backward Euler method is used for the multigroup radiation diffusion, which is treated using multi-grid methods that make use of the HYPRE linear solver library. Currently, since HYPRE does not support OpenMP, parallelism of the multigroup diffusion solver is accomplished through pure MPI.

9.2.2.3 HPC Requirements

In its fullest generality, the radiation transport problem is a seven dimensional problem (three spatial dimensions, one wavelength, one time, and two angles), and therefore has relatively high memory and computational requirements. Calculation of supernova light curves with SEDONA requires particularly high resolution in wavelength if we are to properly resolve the atomic line transitions that provide the primary opacity source for photons. Current SEDONA light curve calculations for 3-D models are typically at a minimally acceptable wavelength resolution (1,000-10,000 wavelength points) and use ~ 128 points per spatial dimension. In the next three years, we aim to achieve much higher wavelength resolution ($10^5 - 10^6$ points) in order to more properly capture the atomic line physics needed to better predict the supernova spectral features and colors. In addition, a higher spatial resolution (~ 512 grid points per dimension) is also desirable to capture small-scale structures (e.g., clumps and fingers) that generally arise due to hydrodynamical instabilities. The current aggregate memory requirements of a SEDONA 3-D light curve (~ 10 GB) thus need to be expanded to ~ 10 TB in the next three years using domain decomposition techniques.

The computational requirements of SEDONA depend largely on the number of radiation pseudo-particles used in the calculation. This value determines the final signal to noise (S/N) ratio of the local grid quantities (e.g., temperature, ionization state) and of the output spectra/light curves. Current typical calculations use 10^8 particles and produce adequate S/N only for integrated quantities (e.g., wavelength or angle averaged quantities). In order to produce useful 3-D output for all wavelengths and viewing angles we will need to increase the particle count to 10^{10} - 10^{11} particles. One advantage of

SEDONA calculations is that they are done in a post-processing mode. Since the hydrodynamical evolution is trivial, only a few time steps (~ 100) are thus needed to resolve the relevant timescales of remnant expansion and photon diffusion.

For core collapse SNe modeling with the CASTRO code, the solution of the radiation transport problem is expected to dominate all other aspects of the explosion simulation (e.g., solving the hydrodynamics or the nuclear reaction networks). CASTRO reduces the dimensionality of the transport problem by adopting multi-group flux-limited diffusion (MGFLD), which neglects angular dependences. High spatial resolution is required (linear resolution of $\sim 10^5$, corresponding to roughly 1000^3 zones for 4-level AMR) while the required wavelength resolution is modest (~ 30 points) because of the smoothness of the neutrino opacities. Because the radiation transport is coupled to hydrodynamics, solution of the multi-group diffusion equation must be performed frequently, e.g., on the dynamical timescale. Many time steps ($\sim 10^4$) are then required for a full 3-D core collapse supernova explosion, thus requiring substantial computational resources.

Previous simulations of supernovae have shown that the outcome of the explosion often depends sensitively on the initial conditions. In observations, we also see that the events span a diverse range of energetics and nucleosynthetic yields. Running a single simulation is therefore not adequate; ensemble calculations are needed to cover a reasonable range of the parameter space (e.g., progenitor star mass and composition, location of ignition conditions). Therefore, we will need computational resources that allow for a significant number (dozen to dozens) of simulations to be carried out.

9.2.2.4 Computational and Storage Requirements Summary

	Current (NERSC 2010)	2014 Requirement
Computational Hours	2 M	50 M
Parallel Concurrency	50,000	500,000
Wall Hours per Run	100k-1M hours	1M-10M hours
Aggregate Memory	10 GB	10 TB
Memory per Core	2 GB	1 GB
I/O per Run Needed	100 GB	500 GB
On-Line Storage Needed	1 TB	10 TB
Archival Storage Needed	5 TB	20 TB
Data Transfer Needed	500 GB	2 TB

9.2.2.5 Support Services and Software

Visualization and data analytics remain a major challenge for astrophysical simulations. For SEDONA light curve calculations, our output consists of ~ 100 GB of data per model, which describes the emergent photon flux as a function of time, wavelength, and viewing angle. For a typical ensemble of calculations, we therefore expect to produce ~ 1 TB of

data, from which we need to extract and visualize the quantities that most effectively probe the physics of the explosion mechanism and nucleosynthesis. In addition, our model data needs to be compared to large and diverse databases of observed supernova spectra in order to evaluate the ultimate success of the simulation. Thus we need useable methods for cross correlating across large databases, which may include irregularly gridded data.

9.2.2.6 Emerging HPC Architectures and Programming Models

The current SEDONA code uses OpenMP threading to distribute the local Monte Carlo particle load, and we are beginning to explore the feasibility of using GPUs to perform similar parallelization of the computational load. There is some hope that Monte Carlo transport methods will be relatively well suited to heterogeneous architectures, as each particle trajectory is an independent thread whose execution time will often be dominated by the computation involved in finding cell boundaries and evaluating scattering kernels. However, data movement will become a significant overhead when particles pass grid boundaries or as the underlying matter distribution evolves with the time step, which will present a challenge for running highly resolved simulations. We are considering several useful approaches, such as scaling up the complexity of the physics treated for each particle (e.g., implementing higher-fidelity scattering kernels) as well as partially computing quantities on the fly (rather than storing them) which should help increase the work load while minimizing data movement. Heterogeneous computing solutions have been shown to be quite powerful for calculating local grid properties – e.g., the thermodynamic state and opacity of each zone – as these problems are often computation heavy.

10 Low Energy Nuclear Theory

10.1 Overview

The Nuclear Theory subprogram provides the theoretical underpinning needed to support the interpretation of a wide range of data obtained from all the other NP subprograms and to advance new ideas and hypotheses that stimulate experimental investigations. This subprogram supports the Institute for Nuclear Theory (INT) at the University of Washington, where leading nuclear theorists are assembled from across the nation to focus on key frontier areas in nuclear physics. The subprogram also collects, evaluates, and disseminates nuclear physics data for basic nuclear research and for applied nuclear technologies with its support of the National Nuclear Data Center (NNDC). The extensive nuclear databases produced by this effort are an international resource consisting of carefully organized scientific information gathered over 50 years of low-energy nuclear physics research worldwide.

10.2 Case Studies

10.2.1 Monte Carlo Calculations of Nuclei and Nucleonic Matter

Principal Investigator: J. Carlson (LANL)

Contributors: S. C. Pieper and R. B. Wiringa (ANL)
S. Gandolfi (LANL)

NERSC Repository: m308

10.2.1.1 Summary and Scientific Objectives

Atomic nuclei are intriguing quantum systems that are important to understand in themselves and for their use in diverse physical processes ranging from stellar evolution and the plasma environment at National Ignition Facility to the fundamental properties of neutrinos. Our ultimate goal is to understand the properties of nuclei from interactions between the nucleons, and to use this understanding to be able to predict as yet unmeasured properties and reactions of nuclei.

In many cases we would like to understand the properties of nuclei as yet unmeasured. For example, an important goal of the Facility for Rare Isotope Beams (FRIB) being constructed at Michigan State University is the ability to study the properties of very neutron-rich nuclei. These nuclei are typically calculated within density functional theory; microscopic theory of nucleonic matter can be used to place severe constraints on

the nuclear density functionals. Such constraints are particularly important when using these theories to go well beyond the regime where they are experimentally tested, for example in neutron stars and supernovae.

Reactions involving the lightest nuclei are also important. The fusion reactions of the lightest elements in big-bang nucleosynthesis form the dominant nuclear species. While many of these reactions are well measured experimentally, others are far more uncertain. Reactions involving unstable nuclei are particularly difficult to address experimentally, and theoretical/computational approaches are particularly important. Reactions with specific energy thresholds can be used, for example, to diagnose the properties of the plasma in NIF capsules.

Nuclei are also critical for understanding physics beyond the standard model. For example, many present-day neutrino experiments use ~ 1 GeV neutrinos produced at accelerators, and it is critical to understand their cross-section with nuclei in order to examine the mixing angles in the neutrino matrix. Neutrinoless double beta decay experiments, where searches for extremely rare processes are undertaken, are critical to study the nature of the neutrinos (Majorana or Dirac) and the overall order and scale of neutrino masses. Accurate nuclear physics calculations are required to calculate the matrix elements relevant for the decay, these matrix elements allow one us to go from a measured rate to the overall scale for the neutrino mass.

10.2.1.2 Methods of Solution

Ab initio nuclear structure calculations require the solution of a strongly correlated quantum mechanical many-body problem. Nuclei are intriguing in that the nucleon-nucleon interaction couples the spin- and spatial degrees of freedom in an essential way. A simultaneous description of both these short-range correlations and the long-range correlations associated with clustering and pairing are required for an accurate description of nuclei.

We employ Quantum Monte Carlo techniques to solve these nuclear physics problems. Complementary methods such as Configuration Interaction (CI) shell model and coupled-cluster techniques have different strengths and combined can offer important insights into the structure and reactions of atomic nuclei. We use different Monte-Carlo approaches to study atomic nuclei: in Variational Monte Carlo (VMC) a parameterized form of the nuclear wave function is assumed, and Monte Carlo methods are used to perform the quadratures necessary to optimize the parameters in the wave function. Green's function Monte Carlo solves for the ground state or low-lying excited states by performing a projection in imaginary time: the ground state is obtained through a path-integral projection to the zero-temperature ground state. Again Monte Carlo methods are used to sample the spatial variables in the path integral. A newer technique, Auxiliary-Field Diffusion Monte Carlo (AFDMC), uses Monte Carlo in both the spin and the spatial degrees of freedom. AFDMC can treat larger systems of up to order 100 particles, though

the accuracy for small systems is not quite as good as GFMC. VMC, GFMC, and AFDMC can treat two- and three-nucleon interactions.

In both the GFMC and AFDMC codes, the desired states are obtained through a branching random walk process. Each step of the random walk involves significant linear algebra. In GFMC this takes the form of large sparse matrix multiplications, the spin alone requires 2^N degrees of freedom, and the wave functions are vectors in a combined spin-isospin basis. In AFDMC dense linear algebra (matrix multiplication and diagonalization) are required at each step. The matrix sizes in AFDMC are typically $9N$, where N is the number of nucleons.

The different random walks can easily be split among the various cores of a large computer; the branching, however, requires significant load balancing between the cores. In GFMC, the linear algebra is large enough to require splitting the linear algebra between cores on a single node, while in AFDMC this has not yet been required.

Parallelism in all our codes is expressed using a combination of MPI and OpenMP. We have been working with Rusty Lusk at Argonne National Laboratory on a new library ADLB (Asynchronous Dynamic Load Balancing) to make a much more efficient, and more general purpose, load-balancing scheme for Green's function Monte Carlo.

10.2.1.3 HPC Requirements

In order to evaluate the path integrals with sufficient accuracy, we must sample many paths over many time steps. Typical calculations involve 25,000 to 500,000 paths, depending upon the size of the system and the accuracy required. Each step in the path evolves the system over a short imaginary time ($\sim 5 \times 10^{-4} \text{ MeV}^{-1}$) that depends upon the structure of the nuclear interaction. To project to a ground or low-lying state, we quite often project to a total imaginary time of order 10 MeV^{-1} , which requires on the order of 20,000 steps for each path.

The load-balancing requirements are significant for the largest calculations. At present GFMC can be used to solve for nuclei as large as Carbon-12. Through the use of OpenMP for the linear algebra within a node, and MPI / ADLB for communication across nodes, we have reached efficiency of up to 85 per cent for 132,000 cores. We are continuing our work with Lusk and collaborators to further increase the efficiency of the GFMC code.

Both the AFDMC and GFMC codes implement checkpoint/restart. We typically write out the present status of all paths at intervals of approximately 2 hours; this information is sufficient to restart the calculation from that point. The I/O requirements are not tremendously large as long as checkpointing is kept to a moderate interval. Sometimes we also use these stored configurations for later analysis. In this case storage requirements can begin to be significant, up to several terabytes per run.

By 2014 we will need significant increases in computational resources, chiefly overall floating point throughput. Advanced applications including calculations of scattering and response functions will require nearly an order of magnitude increase in computer resources over a ground-state calculation for the same nucleus. We will also require significant increases in memory per node. Using OpenMP we have some flexibility in the memory required per core, but the minimum memory per node grows substantially with nucleus size. Ideally we would like to move toward calculations of Carbon-14 by 2014.

Key science tasks going forward include:

- calculations of neutrino response which requires implementing many sets of states along each path;
- calculations of low-energy scattering with multiple channels which will also benefit from multiple sets of stat;
- algorithmic improvements to reduce the communication and memory footprint, in particularly implementing ADLB in AFDMC and improving the implementation and coordination with ADLB generally.

10.2.1.4 Computational and Storage Requirements Summary

	Current (NERSC 2010)	2014 Requirement
Computational Hours	4.7 M	50 M
Parallel Concurrency	Up to 130,000	500,000 or more (at 2-8 GB/node)
Wall Hours per Run	Up to 12 hours	3 to 12 hours
Aggregate Memory	Can use all available memory	Can use all available memory
Memory per Core	2 GB (1.3 GB on Hopper)	2 GB (or more)
I/O per Run Needed	20 GB output 1 TB checkpointing/ configuration saving	100 GB output 10 TB for checkpointing / configuration saving
On-Line Storage Needed	1 TB	20 TB
Archival Storage		
Data Transfer Needed	10 TB	100 TB

10.2.1.5 Support Services and Software

We need BLAS, LAPACK, and SLAPACK for handling dense kernels. We also need PARPACK for some of the eigenvalue calculations

10.2.1.6 Emerging HPC Architectures and Programming Models

We currently use a hybrid OpenMP/MPI programming model for our main production code. We will continue to pursue this model in the immediate future for our production codes since it continues to serve us reasonably well. We will continue to work with the ADLB library and expected continued improvement in scaling up to at least 500K cores as such machines become available.

We have undertaken some experiments with GPU technology. This could possibly be advantageous in the future for AFDMC, where dense linear algebra is required.

Algorithmic changes are expected, but will not likely change the overall structure and requirements of the code.

10.2.1.7 Selected References and Bibliography

“Cold Neutrons Trapped in External Fields,” Gandolfi S.; Carlson J.; Pieper Steven C.; PRL, 012501, 2011.

“BEC-BCS crossover and universal relations in unitary Fermi gases,” Gandolfi S.; Schmidt K. E.; Carlson J.; PRA, 041601, 2011.

“Low-density neutron matter,” Gezerlis Alexandros; Carlson J.; PRC 81, 025802, 2010.

“More Scalability, Less Pain,” E. Lusk, S.C. Pieper and R. Butler, SciDAC Review 17, 30, 2010.

“Strongly paired fermions: cold atoms and neutron matter,” Gezerlis Alexandros; Carlson J.; PRC 77, 032801, 2008.

“Quantum Monte Carlo calculations of neutron-alpha scattering,” Nollett Kenneth M.; Pieper Steven C.; Wiringa R. B.; et al.; PRL 99,022502, 2007.

“Tensor forces and the ground-state structure of nuclei,” Schiavilla R.; Wiringa R. B.; Pieper Steven C.; et al.; PRL 98, 132501, 2007.

“Quantum Monte Carlo calculations of light nuclei,” Pieper SC; Wiringa RB; Annual Reviews of Nuclear and Particle Physics 51, 53, 2001.

10.2.2 *Ab initio* Nuclear Structure Calculations

Principal Investigator: James Vary (Iowa State University)

Contributors: Pieter Maris (Iowa State University)

NERSC Repository: m94

10.2.2.1 Summary and Scientific Objectives

Atomic nuclei are quantum systems that comprise 99.9% of the mass of the visible universe; yet, unlike the situation for atomic electrons, we lack precise knowledge of most of their properties. That is, we do not have a detailed microscopic understanding, based on first principles, of how their constituents, neutrons and protons, are bound together, how the nuclear modes of excitation are formed, or how nuclear reactions take place. This lack of precise knowledge limits our ability to efficiently use atomic nuclei for nuclear energy both for next-generation fission reactors and for fusion reactors under development.

There are reactions that take place in present-day experimental fusion devices and in nuclear reactors that involve unstable nuclei. These reactions are difficult or impossible to measure in laboratory experiments yet they may play important roles in these systems. Take, for example, neutron scattering on the rare nuclei, Oxygen-17 (0.038% natural abundance) and Oxygen-18 (0.2% natural abundance), that are produced in nuclear reactors. These reactions modify the neutron flux in ways that are currently treated with phenomenological models that have unquantified errors. We need a predictive approach for these reactions and many others with quantified uncertainties to help improve the design parameters of nuclear reactors. Our long-term goal is to predict the structure and reactions of rare or even unstable nuclei, and the neutron + Oxygen-17 and the neutron + Oxygen-18 scattering problems are leading examples among many.

Another driver for *ab initio* simulations of atomic nuclei is the quest to solve fundamental problems that will help uncover new laws of nature. A primary example is the search for the exotic process called neutrinoless double beta decay in nuclei, which violates one of the fundamental conservation laws accepted up to the present time. In this process, yet to be seen unambiguously, the nuclear reaction produces a net number of leptons violating the fundamental law of lepton number conservation. There are vigorous worldwide efforts by many groups searching for this rare process in experiments costing tens of millions of dollars. Once this rare process is detected, the nuclear structure physics for the nuclei participating in the decay process will need to be calculated with high precision in order to determine exactly the new laws that violate this conservation law. This research area is dubbed "new physics," or "physics beyond the standard model," and it has major implications for our knowledge of the birth of the universe and the asymmetry between matter and antimatter in the universe. It is acknowledged to be one of the leading science questions for the 21st century (see NAS Report).

10.2.2.2 Methods of Solution

In *ab initio* nuclear structure we deal with a strongly correlated non-relativistic quantum many-body problem. The challenge is to solve the many-body Schroedinger equation for the nuclear wave function, which describes all nucleons inside the nucleus. We treat the nucleons as independent particles, all interacting with each other via state-of-the-art 2-body and 3-body interactions. In order to solve for the wavefunction, we expand it in a finite-dimensional basis of Slater-determinants of single-nucleon states, express the Hamiltonian in this basis, and solve for the lowest eigenstates.

Mathematically, our approach casts the many-body Schroedinger equation as a matrix eigenvalue problem. The computational challenge is to construct the Hamiltonian matrix from the input 2- and 3-nucleon interactions, and to solve it for the lowest 10 to 20 eigenvalues and eigenvectors. (For certain applications we need hundreds of eigenvalues and eigenvectors.) We use the Lanczos algorithm to find the lowest eigenvalues. At the end of the run, we use the eigenvectors to calculate a set of observables to compare with available experimental data or to make needed predictions.

The calculations described above are performed using a suite of codes called MFDn (Many Fermion Dynamics – nuclear).

Parallelism in all our codes is expressed using a combination of MPI and OpenMP.

10.2.2.3 HPC Requirements

In order to obtain accurate solutions of the quantum many-body problem, we have to use very large basis spaces for the expansion of the wave functions. Currently we are using dimensions ranging from 100,000 to several billion. The corresponding Hamiltonian matrix is a very sparse real symmetric matrix. The sparsity depends on the nucleus and on the type of the interaction: with a 3-body interaction the number of nonzero matrix elements is about 10 to 100 times larger than with a 2-body interaction (for the same nucleus and basis space). There is evidence that 4-body interactions will be needed for some nuclei. We anticipate that with 4-body interactions the number of nonzero matrix elements increases by another order of magnitude. In order to converge the 10 lowest states we typically need of the order of 500 Lanczos iterations.

We use a 2-dimensional distribution of the matrix over all available processors. We have a very good load balancing method, in the sense that the nonzero matrix elements are distributed evenly among all processors.

The Lanczos vectors are stored in memory, distributed over all processors, for fast orthogonalization and normalization after every iteration. Again, this is well load-balanced.

Both the storage of the matrix, and the storage of the Lanczos vectors, require a significant amount of memory per core. With our current implementation, based on an even (load-balanced) distribution of the matrix and of the Lanczos vectors over all processors, we get reasonably good performance with 2 GB/core, but our performance degrades if the available memory per core is reduced. That is, as the available memory/core decreases we require more cores for the same calculation and this means increased communications at each Lanczos iteration.

We currently do not have a checkpoint/restart capability, but plan to write each Lanczos vector as it is generated, in order to allow for restarts. This will greatly increase our I/O and on-line storage needs.

By 2014 we will need a better handle of the memory and communication requirements as we continue our calculations with 3-body interactions for larger nuclei, and possibly 4-body forces. We will need algorithmic improvements to reduce the amount of communication and/or to overlap communication with computation in our application. We will need an efficient and load-balanced “on-the-fly” algorithm (in which part or all of the Hamiltonian matrix is not stored but recalculated as needed) in order to reduce our memory footprint.

Key science tasks going forward include:

- develop suitable 4-body interactions, and fit its parameters to data;
- develop and implement a compact representation for 3-body and 4-body interactions;
- algorithmic improvements to reduce the communication and memory footprint.

10.2.2.4 Computational and Storage Requirements Summary

	Current (NERSC 2010)	2014 Requirement
Computational Hours	3.4 M	50 M
Parallel Concurrency	Up to 50k	250k or more (at 2 GB/core)
Wall Hours per Run	Up to 12 hours	3 to 12 hours
Aggregate Memory	Can use all available memory	Can use all available memory
Memory per Core	2 GB (1.3 GB on Hopper)	2 GB (or more)
I/O per Run Needed	40 GB problem input 100 GB output (no checkpointing)	100 GB problem input 0.5 TB output 20 TB for checkpointing
On-Line Storage Needed	1 TB	20 TB
Archival Storage	10 TB	100 TB

10.2.2.5 Support Services and Software

We need BLAS, LAPACK, and SLAPACK for handling dense kernels. We also need PARPACK for some of the eigenvalue calculations

10.2.2.6 Emerging HPC Architectures and Programming Models

We currently use a hybrid OpenMP/MPI programming model for our main production code. Auxiliary pre- and post-processing codes are currently either serial, or MPI only codes. We plan to use a hybrid programming model in order to improve the performance of these pre- and post-processing codes.

Algorithmic changes are expected, both in the pre- and post-processing codes as well as in our main production code, including possibly the incorporation of dynamic load balancing techniques, when/if the number of cores increases on a node. In addition we need to explore the possibility to use GPUs for certain parts of our code.

10.2.2.7 Selected References and Bibliography

National Academy of Sciences National Research Council Committee on Physics of the Universe report, Connecting Quarks with the Cosmos: 11 Science Questions for the New Century (2003); ISBN-13: 978-0-309-07406-3

P. Sternberg, E.G. Ng, C. Yang, P. Maris, J.P. Vary, M. Sosonkina and H. V. Le, "Accelerating Configuration Interaction Calculations for Nuclear Structure", In Proceedings of the 2008 ACM/IEEE Conference on Supercomputing (Austin, Texas, November 15 - 21, 2008). Conference on High Performance Networking and Computing. IEEE Press, Piscataway, NJ, 1-12.; DOI= <http://doi.acm.org/10.1145/1413370.1413386>

M. Sosonkina, A. Sharda, A. Negoita and J.P. Vary, "Integration of ab initio nuclear physics calculations with optimization techniques", Lecture Notes in Computer Science , Vol. 5101, Bubak, M.; Albada, G.D.v.; Dongarra, J.; Sloot, P.M.A. (Eds.) pp. 833 – 842 (2008).

J. P. Vary, H. Honkanen, J. Li, P. Maris, S. J. Brodsky, P. Sternberg, E. G. Ng and C. Yang, "Recent Progress in Hamiltonian light-front QCD," in Light Cone 2008 Relativistic Nuclear and Particle Physics, Proceedings of Science LC2008 Proceedings of Science LC2008 (Sissa, IT), 040 (2008), 040 (2008) ; SLAC-PUB-13460; arXiv: 0812.1819.

N. Laghave, M. Sosonkina, P. Maris and J.P. Vary, "Benefits of Parallel I/O in Ab Initio Nuclear Physics Calculations," (ICCS) 2009 , Part I, (Springer Verlag, Berlin), G. Allen et al., (Eds.), LNCS 5544, 84 (2009).

J.P. Vary, P. Maris, E. Ng, C. Yang and M. Sosonkina, “Ab initio nuclear structure – the large sparse matrix eigenvalue problem,” *Journal of Physics: Conference Series* 180, 012083 (2009). doi:10.1088/1742-6596/180/1/012083; arXiv nucl-th 0907.0209

J. P. Vary, H. Honkanen, Jun Li, P. Maris, S. J. Brodsky, A. Harindranath, G. F. de Teramond, P. Sternberg, E. G. Ng, C. Yang, “Hamiltonian light-front field theory in a basis function approach”, *Phys. Rev. C* 81, 035205 (2010); arXiv nucl-th 0905.1411; DOI: 10.1103/PhysRevC.81.035205

J. P. Vary, H. Honkanen, Jun Li, P. Maris, S. J. Brodsky, A. Harindranath, G. F. de Teramond, P. Sternberg, E. G. Ng, C. Yang, “Hamiltonian light-front field theory within an AdS/QCD basis”, *Nucl. Phys. B* 199, 64 (2010) DOI: 10.1016/j.nuclphysbps.2010.02.008.

P. Maris, M. Sosonkina, J. P. Vary, E. G. Ng and C. Yang, “Scaling of ab-initio nuclear physics calculations on multicore computer architectures,” *International Conference on Computer Science, ICCS 2010, Procedia Computer Science* 1, 97 (2010).

J. P. Vary, H. Honkanen, Jun Li, P. Maris, A. M. Shirokov, S. J. Brodsky, A. Harindranath, G. F. de Teramond, P. Sternberg, E. G. Ng, C. Yang, “Ab initio Hamiltonian approach to light nuclei and to quantum field theory”, *Pramana Jnl of Phys.* 75, 39 (2010) DOI 10.1007/s12043-010-0063-2

F. Liu, R. Mundhe, M. Sosonkina, C. Cockrell, M. Aronnax, P. Maris and J. P. Vary, “A Data Management System for Ab-Initio Nuclear Physics Applications,” accepted for publication in the proceedings of High Performance Computing – 2011 (HPC-2011)

H. M. Aktulga, C. Yang, E. G. Ng, P. Maris and J. P. Vary, “Large-scale Parallel Null Space Calculation for Nuclear Configuration Interaction,” accepted for publication in HPCS-2011.

A. Srinivasa, M. Sosonkina, P. Maris, J. P. Vary, “Dynamic adaptations in ab-initio nuclear physics calculations on multicore computer architectures,” *Proceedings of 12th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC-11)* (to appear).

10.2.3 Nuclear Density Functional Theory

Principal Investigator: Jonathan Engel (University of North Carolina)

Contributors: Joseph Carlson (LANL)

NERSC Repository: m308

10.2.3.1 Summary and Scientific Objectives

Nuclear Density-Functional Theory (DFT) refers to calculations that incorporate some kind of mean-field equations. The approach is based on an energy-density functional and can include effects (correlations) that are beyond usual mean-field theory but have been repackaged into mean-field-like equations or extensions of such equations that superpose mean-fields. The energy functionals in use at present are still largely phenomenological — that is, not derived directly from real (measured) interactions between two and three nucleons. For this reason *ab initio* techniques such as Green’s Function Monte Carlo or the No-Core Shell Model, which do make direct use of the real interactions, are probably better than DFT in light nuclei. In heavy nuclei, however, *ab initio* methods (and even non-*ab-initio* approaches such as the shell model with an excluded core) become impossible to apply and DFT, despite its still phenomenological character, is usually the method of choice.

A significant part of the nuclear structure community works in the DFT framework. Here we describe research by two groups, one on static properties of heavy nuclei and one on dynamic processes such as collisions and fusion. The two programs together make up much of the computationally intensive work in the field.

Static properties: The goal here is to precisely calculate nuclear masses, spectra, beta-decay rates, fission probabilities, etc., both for their own sake and to help understand particles such as neutrinos that act inside nuclei or astrophysical objects such as neutron stars, which are essentially giant nuclei.

Dynamic properties: The approach discussed here is called the time-dependent superfluid local density approximation (TD-SLDA). The method is designed to help understand heavy-ion collisions, large-amplitude nuclear oscillations, fusion, vortices in the crusts of neutron stars, and even dynamics in nonnuclear systems such as cold atomic gases.

10.2.3.2 Methods of Solution

Static properties: The code HFODD is the central piece of the effort. It generates many mean-field configurations, which are combined to compute properties of nuclear states. HFODD is made of a kernel and a MPI interface. The kernel represents the Hartree-Fock-

Bogoliubov (HFB) equations, a non-linear system of coupled integro-differential equations, as matrix equations in a harmonic-oscillator basis. The equations thus take the form of a pseudo-eigenvalue problem in which the matrix elements of a Hermitian matrix depend on the eigenfunctions themselves. The resulting fixed-point problem is solved iteratively with standard acceleration techniques (Anderson and Broyden methods, etc.).

The code uses hybrid MPI/OpenMP to simultaneously compute a large number of different nuclear configurations, each of which is characterized by a set of constraints on the particle number, nuclear shape, angular momentum, etc. The MPI interface coordinates the calculation of these configurations, currently in a round-robin fashion. Individual configurations are handled through a single MPI process by the kernel, and are accelerated with OpenMP instructions. The procedure for diagonalizing the always-dense Hermitian matrices uses BLAS and LAPACK libraries.

Dynamic properties: The TD-SLDA produces a set of time-dependent mean-field (Bogoliubov de Gennes-like) equations — complex nonlinear partial differential equations with periodic boundary conditions. Typically, a static problem must be solved first to produce an initial condition; the methods used for that are similar to, though not identical with, those associated HFODD and described above. One difference is that here the wave functions are represented on a cubic lattice rather than an oscillator basis.

In evolving the solutions forward in time, The TD-SLDA method evaluates the action of the Hamiltonian on the wave functions, so the determination of eigenvalues is not necessary. To speed up the evaluation of first- and second-order derivatives, the FFTW library is used, so that matrix operations are avoided altogether. This allows the evaluation of spatial derivatives with extremely high accuracy and with essentially the same speed as in a multi-step finite difference formula. The time evolution itself is performed with a multistep, fifth-order predictor-corrector-modifier Adams-Bashforth-Milne method. The method requires only two evaluations of the right hand side of the differential equations per time step, a number that cannot be reduced without going to a lower-accuracy numerical method. The timestep is chosen so that the relative truncation error in the Adams-Bashforth-Milne method is between 10^7 and 10^{15} . Simulation results for runs with up to a few million time steps show that the method is stable and accurate.

In each time dependent simulation various time dependent quantities are stored on the lattice for use in later studies.

10.2.3.3 HPC Requirements

Static properties: About 100 mean-field configurations are currently used for such typical problems as understanding fission or doing global surveys of nuclear properties. For each configuration the size of the matrix dimensions are of order 2,000, and the number of iterations needed to reach convergence varies between about 50 (ground-state) and 5,000 (for extreme deformation of the kind produced by fission.) Typical jobs process about 10,000 configurations at a time and hence require about 10,000 MPI processes with 6

threads/process. We find that 150 HFODD iterations correspond approximately to ten hours of wallclock time on Hopper.

The performance of the kernel is the major bottleneck. Solving the HFB equations in a large basis is computationally intensive because of the large number of dense-matrix operations. The precision of the calculation is directly proportional to the number N of oscillator states in the basis. The HFODD kernel, however, currently operates on one MPI process only, and the calculation time therefore approximately scales as N^8 .

By 2014, the HFODD kernel will be spread across several MPI processes in order to (i) increase the precision of the through larger bases, (ii) explore richer energy-density functionals (for example functionals of the non-local density matrix), and (iii) implement more advanced and load-balanced task-management systems, which, in turn, will approximately decrease the total number of configurations to be processed.

Dynamic properties: As of now, the codes routinely use over 97% of all of JaguarPF To calibrate demand and for economy of space, the table below lists raw time-step data averaged over hundreds of steps that include observable write events at regular frequencies (as measured on JaguarPF in production-scale problems):

Code	Lattice	T[s] / ts	FP OP / ts	Cores
Atomic gas	$50^2 \cdot 100$	6.30	1.58E+14	103,917
Nuclear	$40^2 \cdot 64$	6.93	9.45E+13	136,628

Each process evolves a single wave function. In the nuclear case there are two additional processes to coordinate various computations between proton and neutron “working groups.” In all, the atomic-gas system equates to evolving 207,834 coupled nonlinear PDEs and the nuclear system evolves 546,504 such equations. Given the data, 100,000 time steps with this realistic atomic-gas scenario will cost 6.3 (s/time step) \cdot 100,000 (time steps) \cdot 103,917 PEs / 3,600(s/hr) \approx 18M CPU hours. Thus, 100,000 time steps of the nuclear code for the lattice in the table will thus cost about 26M CPU hours.

Today, the nuclear code can be used to simulate the excitation of nuclear resonances by a projectile such as an electron. A single run requires between 15,000 and 20,000 time steps per orientation and impact parameter. In total we will execute from 20 to 30 runs to study the combinations. The nuclear solvers cost roughly 2M CPU hours for this size system. Thus, we estimate 300,000 time steps \cdot 26M CPU hours / 100,000 time steps + 2M CPU hours \approx 80M CPU hours. A shock-wave/domain-wall study with the atomic-gas codes will require 18M CPU hours. About 1M CPU hours are spent in checkpoint/restart. Thus, 100M CPU hours are needed for problems that one should attack this year.

For next year, the induced fission problem will require long run times (100,000 to 150,000 time steps) on a possibly more complex lattice. There are different ways to induce fission and each will require its own studies. By 2014 the plan is to superpose many time-dependent mean-field configurations (through a process called stochastic

quantization). To avoid the need for exascale facilities, pairing will be neglected initially, simplifying the underlying mean-field problem. That should keep usage to about 100M CPU hours, but ideally much more would be available. If more advanced/larger scale hardware emerges, pairing will be included, on small lattices initially.

10.2.3.4 Computational and Storage Requirements Summary

Static Properties		
	Current (NERSC 2010)	2014 Requirement
Computational Hours	4.7 M	5 M
Parallel Concurrency	10K MPI process	100K MPI processes
Wall Hours per Run	10	12
Aggregate Memory	120 TB	1200 TB
Memory per Core	2 GB (with SMP size=24 MB)	1 GB (with same SMP size)
I/O per Run Needed	400 GB	4 TB (or more)
On-Line Storage Needed	1 TB	10 TB
Archival Storage	100 TB	1,000 TB
Data Transfer Needed	Negligible	Negligible

Dynamic Properties		
	Current (NERSC 2010)	2014 Requirement
Computational Hours	4.7 M	100 M
Parallel Concurrency	10K to 153K PEs	Same
Wall Hours per Run	6 to 175	Same
Aggregate Memory	20TB to 212 TB	Same
Memory per Core	$O(2^{31})$ B	Same
I/O per Run Needed	$O(2^{43})$ B on avg (cprs)	Same
On-Line Storage Needed		
Archival Storage	$O(2^{45})$ B	$O(2^{50})$ B
Data Transfer Needed	$O(2^{31})$ B/run	same

10.2.3.5 Support Services and Software

Static properties: The problem of finding an efficient and scalable parallel eigensolver for dense Hermitian matrices is relatively general, and is currently included as a research topic in the FASTMath SciDAC-3 Institute. The effort described here would greatly benefit from interactions with the people involved. Code profiling can be performed much more efficiently by experts than by physicists and the project could use help there. Finally, help in identifying parts of the HFODD kernel that could benefit from GPUs would be much appreciated.

Dynamic properties: The (TD-)SLDA software is written primarily in C with support drivers in Fortran. Existing software libraries are used whenever possible. Here is a short

list of typical library dependencies: MPI, Cray's Scientific Library (LibSci for BLAS, BLACS, ScaLAPACK); POSIX Threads (pthreads); FFTW; and LUSTRE (including liblut). VisIt is used for visualization and the Silo library for data. Apple's iMovie is used for production quality movies of simulation results. The SLDA software employs a scalable (Lustre-based) high performance checkpoint and restart scheme for simulation continuity in both the static and time-evolution codes.

10.2.3.6 Emerging HPC Architectures and Programming Models

Static properties: In the past 5 years, HFODD has evolved from a 20 year-old serial legacy code (1997-2007) to an MPI-only code (2007-2009) to a hybrid MPI/OpenMP programming model (from 2009). For the next 3 years, the plan is to implement a Hybrid MPI/OpenMP model on the kernel itself, and explore how some of the computationally intensive vectorized routines in the kernel could benefit from GPU architectures.

Dynamic properties: The time-evolution algorithm can be further parallelized in a coarse manner by separating the individual components of the wave functions. Fine-grain opportunities exist in individual terms that appear in function evaluation or in overlapping different tasks related to I/O and communication with processing. It may also be that the individual wave function components are better handled with lightweight processes. Transporting the codes to hybrid hardware platforms requires rewrites of portions of the software. OpenCL, CUDA, and FUSION compiler techniques are being investigated at UW, on DIRAC at NERSC, and (soon) on Keeneland at ORNL. A new postdoctoral researcher with expertise in coding GPUs for lattice QCD is working with Roche to investigate the interplay of multi-core hosts with GPU accelerators.

11 Heavy Ion Experiments

11.1 Overview

A major component of the U.S. program in Nuclear Physics addresses fundamental issues related to the role of partons — quarks and gluons — in the character of the nuclear force, the properties of nuclear matter at both normal and high energy density, and the origins of the universe as it evolved after the big bang. These investigations lead to experimental tests of the theory of quantum chromo-dynamics (QCD), the mechanisms of quark confinement and chiral symmetry, the search for new forms of matter such as the quark-gluon plasma, and confirmation at modern accelerators of the QCD phase transition in the early universe.

The Heavy Ion subprogram investigates the high temperature frontier of QCD, by trying to recreate and characterize new and predicted forms of matter and other new phenomena that might occur in extremely hot, dense nuclear matter and which have not existed since the Big Bang. Measurements are carried out primarily using relativistic heavy ion collisions at RHIC, the Relativistic Heavy Ion Collider at Brookhaven National Lab. Participation in the heavy ion program at the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN) provides U.S. researchers the opportunity to search for new states of matter under substantially different initial conditions than those provided by RHIC, yet still provide information regarding the matter that existed during the infant universe.

11.2 Case Studies

11.2.1 The ALICE and STAR Heavy Ion Collider Experiments

Principal Investigator: R Jeff Porter (LBNL)

Contributors: Richard Jones (U. of Connecticut and the GluEx experiment)

NERSC Repositories: ALICE and STAR

11.2.1.1 Summary and Scientific Objectives

The science program at the U.S. Relativistic Heavy Ion Collider (RHIC) facility has provided pioneering discoveries of a new state of thermalized matter, known as the Quark Gluon Plasma (QGP). The STAR experiment (Solenoidal Tracker at RHIC), one of the four original RHIC experiments, has been a leader in such discoveries, exploring the properties of the QGP and illuminating aspects of this new state that were surprisingly

different than theoretical predictions made prior to the RHIC era. The STAR program continues at RHIC providing new measurement from new detectors, higher statistics and at different beam energies in order to further map out the properties of the matter and determine threshold energy densities for the onset of QGP formation.

The ALICE (A Large Ion Collider Experiment) collaboration constructed and operates a heavy-ion detector to exploit the unique physics potential of proton-proton and nucleus-nucleus interactions at the even larger Large Hadron Collider (LHC) energies. The science program with ALICE extends the work done at RHIC with a principal goal of studying the physics of this new phase of strongly interacting matter at extreme energy densities.

The physics of the QGP is studied by exploring the properties of heavy ion collisions using a multitude of signals within each collision. These signals include global event properties such as particle multiplicities and event-type characterizations, identified particle distributions in transverse energy and momentum, particle correlations, and rare or heavy particle yields. As a result, the science is conducted by hundreds of physicists working on dozens of investigations from common event samples to explore different aspects of the collisions and infer information about the QGP and its formation.

11.2.1.2 Methods of Solution

In collider-based experiments like STAR and ALICE, collisions are registered within the detector as independent events. A computing task is the processing of an event collection, done as a set of independent jobs with each job assigned a subset of the collection. This type of processing is “pleasantly parallel” and easily distributed onto clusters of off-the-shelf hardware. Many of the specific investigations that the scientists pursue require very large data samples to make statistically significant measurements. As such both STAR and ALICE have yearly data acquisition targets approaching 10^9 events, corresponding to Petabyte scale annual data volumes.

Event processing is carried out in several stages. Initial pattern recognition and detector calibration tasks are done at the experiment site on dedicated compute facilities. Once the calibrations are finalized, a small team of scientists run the data through a production process known as reconstruction, which takes the experiment data from its raw form of detector signals and re-assembles each event as a reduced set of interpretable physics quantities such as particle tracks and energy deposition in calorimeters. These reconstruction tasks run over hundreds of TBs of data requiring several weeks or months of processing on large, sometimes distributed, compute farms. As a result, new event summary files are produced and made available to all physicists in the collaboration for their analysis. Large groups of physicists, however, need only a greatly reduced common subset of information from these event summary files. Thus, formally in ALICE and less formally in STAR, an additional production pass is done over the data to produce a new set of reduced analysis-ready data files in which some investigation-specific pattern recognition algorithms have been applied. Both the event summary files (~100s TBs)

and analysis ready data files (~10s TBs) retain their event-based granularity allowing for further processing to be distributed over many nodes or even many clusters and are made available to individual physicists within the experiment software infrastructure for their analysis tasks.

A parallel set of event processing is also done to understand the efficiencies and systematic effects from the data processing, and is carried out by simulating the collisions and the detector responses to the events. These simulations are then subjected to the same reconstruction process and analysis as done with the real data. The resulting analyses can be compared directly with the known simulated input, giving the scientists the tools to evaluate and quantify the ability of the experiment methods to measure quantities from the collisions. Each investigation topic's analysis treats the event quantities differently and must have the processing effects on their results evaluated separately.

The event-processing tasks described above are generally referred to as production tasks. That is, they are managed by a small number of scientists to produce distilled event files used in analysis. A specific investigation or analysis is itself a processing task but run independently by individual scientists. The goal of such processing is to refine the algorithms and selection criteria used to extract the signal of interest thereby providing the most robust measurement possible, maximizing the signal purity and minimizing systematic biases in the procedures. While there are a variety of techniques used to optimize an analysis, the process is characterized by numerous passes over a large (multi-TB) subset of the data, requiring fast turnaround times (~hours) to evaluate the many facets of the analysis and the physical properties being measured.

The large data sets and large number of physicists accessing the experiment data require use of a common software infrastructure. Both STAR and ALICE have built their infrastructure upon the ROOT data analysis framework, an organized set of C++ class libraries with utilities to handle data I/O, data collection management, histograms and plotting, functional fits to data and many other useful tools. Experiment-specific software is written in C++ and consists of several hundred thousand to more than a million of lines of codes with explicit reliance on the ROOT framework. In addition, simulations used for studying experiment efficiencies and systematic effects rely on the GEANT (or GEANT4) Detector Description and Simulation Toolkit to describe the geometry and materials of the detector and their support structures and to simulate the passage of particles through the material. Maintaining larger infrastructures such as these on multiple platforms can be a challenge. STAR software, for example, is validated on a very limited set of platforms, currently versions of Scientific Linux. The ALICE computing team has been able to focus more on code portability; packaging dependencies within the software releases. ALICE software can and is run on a large variety of modern Linux systems.

The common features of STAR and ALICE workflows – trivially parallel processing of large data volumes on commodity hardware – have led these groups to invest in the NERSC PDSF system. The PDSF model is different than other NERSC systems in that

the cluster hardware is purchased by the scientific groups but operated and located at NERSC to have access to other important NERSC resources. For ALICE and STAR, that important resource is the large capacity archival storage on HPSS. Both experiments use NERSC/HPSS for custodial storage of raw and processed data. In addition, the challenge of maintaining disk residence of all data is mitigated by periodically re-staging large data sets from HPSS to disk. It should also be noted that PDSF manages very large disk capacities. As a point of comparison, PDSF consists of about 2TB disk per core (2PB of disk and 1k cores) while other NERSC systems consist of about 0.05TB/core (10PB of disk and 200k cores).

11.2.1.3 HPC Requirements

The event-processing described in the previous section naturally splits into production and analysis type tasks for the purpose of specifying their resource requirements. Production tasks are characterized by being CPU intensive with I/O requirements of 10-100 KB/s per job and fit well within a 2 GB/core memory allocation. Analysis jobs can be more for I/O- and/or memory-intensive. The majority of analysis tasks are I/O bound which could absorb reading data at 10s MB/sec. There are significant classes of jobs, particularly in ALICE, that also require large memory footprints of 4-5 GB/core. About one-half of the work done on NERSC/PDSF by both STAR and ALICE consists of analysis jobs. Currently neither production nor analysis jobs require or make use of parallel processing techniques.

In order to minimize I/O contention in analysis work, ALICE has adopted running of analysis “trains.” These trains co-locate many analyses tasks into one process, which share reading of input data and accessing the data in memory. A single collaborator, referred to as the conductor, checks out all scheduled analyses from the code management system, compiles the code and runs a test job prior to submission. Once the test job is verified, the conductor runs the train to process a full data set – several hundred independent but identical jobs processing separate subsets of the full data set. A new train is run every (few) day(s) allowing scientists time to evaluate results and refine the algorithms or selection criteria specific to their analysis.

A significant challenge to both ALICE and STAR is the management of the large datasets and their access by the research physicists doing the analysis work. The large numbers of physicists and investigations translate to a large number of independent tasks accessing the same data. The event-based granularity of the data allows for a distributed computing model with data duplicated and distributed onto multiple resources, at computing centers remote to the primary facility of the experiment. Such distributions of data, along with a corresponding distribution of the software infrastructure and user base, expand the scientists’ opportunity to access and analyze the data. However, this requires a level of coordination between operations at all sites in terms of data migration and cataloging, experiment software maintenance, and access by *all* collaborators.

The distribution of data and processing in STAR has generally been limited to the

experiment's primary site on the RHIC/ATLAS Computing Facility (RACF) at BNL and its secondary site at NERSC/PDSF³⁴. STAR users access data and processing resources at PDSF typically by logging directly onto the system; however, STAR does make use of the Open Science Grid (OSG) interface on PDSF for grid submission of some production tasks. The scope of the challenge for STAR changed starting in 2010, when the data volumes dramatically increased over those of previous years. Now new event summary and analysis-ready data files have grown to more than 500 TB per year. Unlike raw data files that are processed at RACF and archived on RACF HPSS, these event summary and analysis-ready data files need to remain disk resident for several years at one or more facilities while the scientists extract the physics measurements. The data files are managed in a global STAR File Catalog and made accessible either directly on networked file systems (NFS or GPFS) or on remote file systems via XRootD⁵, a software framework for fast, low latency and scalable data access. In particular, local disks on compute nodes are made accessible using XRootD, forming a PetaByte-scale disk-based data storage system directly on the cluster.

The distributed processing in ALICE is much more extensive, making almost exclusive use of grid technologies. The ALICE Grid⁶ is made up of almost 80 sites worldwide of which NERSC/PDSF is a member facility. Jobs are run on PDSF via a grid submission model whereby ALICE-specific infrastructure on PDSF pulls work from a central task queue located at CERN. Data are distributed by automated mechanisms onto Grid-enabled ALICE-XRootD based Storage Elements (SE) located at all ALICE Grid sites and managed within the global ALICE File Catalog. ALICE maintains several hundred TBs of disk space on its SE located at PDSF. Jobs are typically run on Grid sites where the data exists and, for most ALICE Grid sites, individual scientists do not require direct logins or have individual user accounts. The ALICE Virtual Organization (VO), in accordance with rules set out by the International Grid Trust Federation⁷ (IGTF), maintains user management separately.

11.2.1.4 Computational and Storage Requirements Summary

³ A few smaller University-based facilities exist; however, burdens of software and user support or network limitations has led many groups to instead provide institutional resources directly on RACF.

⁴ A third major site at KISTI in Korea is being made available to STAR.

⁵ <http://xrootd.slac.stanford.edu/>

⁶ <http://pcalimonitor.cern.ch/map.jsp>

⁷ <http://www.igtf.net/>

	Current (NERSC 2010)	2014 Requirement
Computational Hours	STAR: 5×10^6 ALICE: 3×10^6	STAR: 5×10^6 ALICE: 12×10^6
Parallel Concurrency	100-1000 non-communicating jobs	Same as 2010 or perhaps $\sim 100 \times 10$ -way parallel jobs
Wall Hours per Run	500-5000	500-5000
Aggregate Memory	2-40 TB	2-40 TB
Memory per Core	2-4 GB/core	2-4GB/core
I/O per Run Needed	10-300 TB	10-300 TB
On-Line Storage Needed	0.4 PB/ experiment	1.5 PB/experiment
Archival Storage Needed	0.5 PB/experiment	1.5 PB/experiment
Data Transfer Needed	1Gbps	2-5 Gbps

11.2.1.5 Support Services and Software

The primary support services important to both STAR and ALICE include reliable access to HPSS, and well-maintained network connection to ESnet and data transfer infrastructure supporting WAN data movement. Also, both groups make use of OSG grid interface for job submission onto PDSF. Both groups also request acceptance of production accounts at NERSC so as to simplify production tasks run by a small group of users on behalf of the collaboration. Other support services important to STAR and ALICE differ.

NERSC/PDSF is the principal secondary analysis facility for STAR and has been for more than a decade. Each STAR collaborator is granted login access to PDSF via normal NERSC account procedures and STAR plans to continue running on PDSF in this manner. ALICE, however, is strictly a Grid-based operation consisting of a large number of sites. Thus, requiring each of the several hundred scientists to maintain individual accounts at every site is not a workable solution. As such, *ALICE strongly requests that new user authorization procedures be allowed that rely on the PI and the VO to manage its user base in accordance with common user agreements and practices.*

STAR has determined that Cloud computing based on VM technologies is an extremely good match for its needs. STAR has a significant amount of legacy code that it is unable to port to new operating systems and had been forced to manage its large software infrastructure to support only a small number of systems. VM technology allows STAR to encapsulate these code and OS dependencies in a snapshot environment for processing now and long into the future. STAR has shown it can adapt to a variety of Cloud systems including, Amazon EC2, Eucalyptus on Magellan, and Nimbus at ANL. ALICE, however, has thus far focused on code portability and does not (as of yet anyway) see a general benefit to VM technologies for actual data processing.

ALICE also supports a specific type of analysis facility that is not strictly grid-based. Such facilities rely on PROOF⁸, a parallel ROOT facility framework. Users login into a PROOF facility and request a data set matched with an analysis tasks. The PROOF architecture stages the data with XRootD and optimizes the processing by dynamically feeding data to the worker processes. Advantages of a PROOF facility are the simplicity of configuring an analysis job and the fast overall processing times. There is no PROOF facility at NERSC but both ALICE and GluEx would be interested in exploring the viability of having such a capability at NERSC on PDSF or elsewhere.

11.2.1.6 Emerging HPC Architectures and Programming Models

Because of the natural parallelism in event-based processing, the field has been able meet its ever growing computing needs with faster processors found in commercial off-the-shelf hardware. The community recognizes the need to adapt to and take advantage of changing architectures such as many-core processors. Some independent efforts exist in the field with participation from the experiments; however, neither STAR nor ALICE has an established program to develop general solutions. There exists significant amounts of concurrency in the data and tasks but new prototypes need to be developed to exploit those processing features. Developing and testing such prototypes on NERSC test systems and with NERSC support would be of great value. Fundamentally, the work will require a community effort to produce changes at the code-base levels, ROOT and GEANT.

⁸ <http://root.cern.ch/drupal/content/proof>

12 Nuclear Physics Accelerator Research

12.1 Overview

The Nuclear Physics program supports a broad range of activities aimed at research and development related to the science, engineering, and technology for accelerators of electrons, protons and heavy ions. Research and development is supported that will advance fundamental accelerator technology and its applications to nuclear physics scientific research. Areas of interest include the basic technologies of the Brookhaven National Laboratory's Relativistic Heavy Ion Collider (RHIC), with heavy ion beam energies up to 100 GeV/amu and polarized proton beam energies up to 250 GeV; technologies associated with RHIC luminosity upgrades; the development of an electron-ion collider (EIC); linear accelerators such as the Continuous Electron Beam Accelerator Facility (CEBAF) at the Thomas Jefferson National Accelerator Facility (TJNAF); and development of devices and methods that would be useful in the generation of intense rare isotope beams for the next-generation rare isotope beam accelerator facility (FRIB). A major focus in all of the above areas is superconducting radio frequency (RF) acceleration and its related technologies. Also, as recommended in the last NSAC Long Range Plan, accelerator research and development on the most challenging technical issues related to the EIC, the future Electron Ion Collider, will soon be supported.

12.2 Case Studies

12.2.1 Parallel Simulation of Electron Cooling Physics and Beam Transport

Principal Investigator: D.L. Bruhwiler (Tech-X Corporation)

Contributors: B.T. Schwartz, V.H. Ranjbar, G.I. Bell (Tech-X Corporation)

NERSC Repository: m327

Other m327 Users: J. Qiang (LBNL); S. White and Y. Luo (BNL)

Collaborators: R. Ryne (LBNL); V.N. Litvinenko (BNL) W. Fischer (BNL); G. Wang (BNL); Y. Hao (BNL); K. Paul (Tech-X) and (Tech-X)

12.2.1.1 Summary and Scientific Objectives

This NERSC repository supports NP-funded accelerator modeling activities of the Community Petascale Project for Accelerator Science and Simulation (ComPASS) SciDAC-2 project. This team includes personnel from Brookhaven National Lab (BNL), Thomas Jefferson National Accelerator Facility (JLAB) and Tech-X Corp. The team has been charged to make effective use of emerging petascale hardware to accurately simulate key aspects of the CEBAF upgrade underway at JLAB and various luminosity upgrade possibilities for the Relativistic Heavy Ion Collider (RHIC) at BNL. CEBAF, the

Continuous Electron Beam Accelerator Facility, is the main research accelerator at JLAB. Another primary goal is to reduce technical risk for future Electron-Ion Colliders (EICs) such as eRHIC (a BNL concept for an electron beam at RHIC) and a medium energy electron-ion collider (MEIC, a JLAB concept).

The present focus of parallel simulations at NERSC spans three technical areas:

- electron cooling of relativistic hadron beams (to increase luminosity);
- beam-beam collisions (determine effect on beam dynamics, luminosity);
- spin-tracking (understand how to keep polarized beam fraction high).

In the next three years we expect to

- support the Coherent Electron Cooling (CeC) proof-of-principle experiment underway at BNL;
- provide larger-scale support for near-term RHIC efforts & also EICs;
- provide beam-beam capabilities to support RHIC, LHC, ELIC design.

Coherent Electron Cooling uses a Free Electron Laser (FEL) to combine electron & stochastic cooling concepts. It has three major subsystems:

- Modulator: the ions imprint a “density bump” on the electron distribution;
- Amplifier: FEL interaction amplifies the density bump by orders of magnitude;
- Kicker: the amplified & phase-shifted electron charge distribution is used to correct the velocity offset of the ions.

Electron cooling has been successful at low-energy ion accelerator facilities around the world [1,2,3]. It has been demonstrated [4] that electron cooling can work with a moderately relativistic electron beam. However, the effectiveness of conventional electron cooling systems (i.e. the cooling rate) decreases roughly as γ^{-2} , where γ is the relativistic factor of the ion beam, and this can be larger than 100.

Coherent electron Cooling is an important priority for RHIC & the future Electron-Ion Collider. The 2007 Nuclear Science Advisory Committee (NSAC) Long Range Plan recommends, “...the allocation of resources to develop accelerator and detector technology necessary to lay the foundation for a polarized Electron-Ion Collider.” (See the NSAC website: <http://www.er.doe.gov/np/nsac/index.shtml>.) The 2009 Electron-Ion Collider Advisory Committee (EICAC) selected CeC as one of the highest accelerator R&D priorities. (See the EIC Collaboration website <http://web.mit.edu/eicc>.)

For the proposed Electron Ion Collider the high energy of the ion beam is a serious challenge. Hence, the recent concept of Coherent Electron Cooling⁵ is being explored.

[1] V.V. Parkhomchuk and A.N. Skrinsky, “Electron cooling: 35 years of development,” Rep. Prog. Phys. **54**, 919 (1991).

[2] M. Larsson, “Atomic and molecular physics with ion storage rings,” Rep. Prog. Phys. **58**, 1267 (1995).

[3] I.N. Meshkov, “Electron cooling – the first 30 years and thereafter,” Nucl. Instrum. Methods Phys. Res., Sect. A **391**, 1 (1997).

[4] S. Nagaitsev, D. Broemmelsiek, A. Burov et al., “Experimental Demonstration of Relativistic Electron Cooling,” Phys. Rev. Lett. **96**, 044801 (2006).

[5] V.N. Litvinenko and Y.S. Derbenev, “Coherent electron cooling”, Phys. Rev. Lett. **102**, 114801 (2009).

Many technical difficulties must be resolved via full-scale 3-D simulations before the CeC concept can be validated experimentally.

There are alternative cooling approaches but they are not as effective. For example, stochastic cooling has shown great success with 100 GeV/n Au⁺⁷⁹ in RHIC; however, it will not work with 250 GeV protons in RHIC. High-energy unmagnetized electron cooling could be used for 100 GeV/n Au⁺⁷⁹; however, the cooling rate decreases as $1/g^2$, which is too slow for 250 GeV protons. In contrast, CeC could yield a six-fold luminosity increase for polarized proton collisions in RHIC. This would help in resolving the proton spin puzzle. It could also break the $1/g^2$ scaling of conventional electron cooling, because it does not depend on dynamical friction.

Improvements in NERSC computing hardware, software and services could help us perform faster beam-beam and spin-tracking simulations with greater physical fidelity. This in turn, could provide physical insight that points to beam dynamics changes that significantly increase the luminosity of RHIC with greater polarization and would reduce the time / cost required for obtaining important nuclear physics results and perhaps enable new results. By providing computational support to the CeC proof-of-principle experiment at BNL additional NERSC resources could help that effort succeed, resulting in a fundamentally new and important technique to increase the luminosity of RHIC or of any future EIC facility by orders of magnitude.

The NERSC m327 project repository also supports beam-beam simulations of relevance to the DOE Office of Nuclear Physics. Such simulations, some of which are part of the COMPASS SciDAC project, are important for RHIC and will be of essential importance to the future EIC. Also, repository m327 is not the only repo supporting accelerator technology. Other relevant NP activities in massively parallel accelerator modeling and design include the following:

- SLAC National Accelerator Laboratory
 - FEM modeling of SRF cavities at JLAB, MSU/FRIB
- Lawrence Berkeley National Laboratory
 - parallel particle tracking & beamline design
 - additional beam-beam simulations for MEIC, RHIC and LHC
- Argonne National Laboratory
 - FEM modeling of SRF cavities for FRIB
 - parallel particle tracking & beamline design for FRIB
 - Vlasov/Poisson algorithm development
- Tech-X Corporation
 - Finite-difference time domain (FDTD) modeling of SRF cavities for JLAB
 - inverse cyclotron for light-ion stopping at FRIB
 - electron gun modeling for BNL (diamond amplifier project)

12.2.1.2 Methods of Solution

There are three main code suites used in this work: VORPAL, BeamBeam3D, and UAL-SPINK.

The parallel VORPAL framework employs a variety of algorithms for FDTD simulations of various types of fields and particles. Fields and fluids are represented on a structured Cartesian mesh, while particles move through space. For this project, VORPAL is being used to model non-relativistic Coulomb collisions between electrons and ions in a frame moving with the two distributions. To date, we have primarily used two special-purpose, gridless, N-body algorithms in VORPAL for molecular dynamics simulations of binary Coulomb collisions between the freely scattering charged particles. We have also used a Poisson solver, based on the Trillinos suite of libraries (Aztec, ML), for treating the same problem via electrostatic Particle-in-Cell (PIC). We expect to use the PIC approach more extensively in the coming year, as we move toward much larger domain sizes, because the N-body algorithms do not scale well to large domains. We are also using the delta-f PIC capabilities of VORPAL.

For electrostatic PIC, non-relativistic particle motion in VORPAL is modeled by the 2nd-order leapfrog algorithm due to Boris. High-order spline-based particle shapes are used for charge deposition onto the grid, and in order to conserve momentum, to interpolate forces to the particle locations. This helps to reduce field noise and particle push errors. The gridded charge density is used to solve Poisson's equation for the potential, which is then finite-differenced to obtain E on the standard Yee mesh (see, e.g., the text by Birdsall and Langdon). A semi-analytic N-body algorithm, that we call the “binary Coulomb collision” method (BCC), is able to accurately capture strong electron/ion scattering events in the presence of arbitrary external fields with global 2nd-order accuracy [6].

For electrostatic PIC, the Trillinos library suite is used to iteratively solve the Poisson equation in parallel. For the N-body BCC algorithm, analytic two-body formulae (see e.g., the standard text by Goldstein) are used for close interactions, while a 2nd-order accurate expansion of these equations is used for weak/distant interactions. External fields are included with a standard leap-frog based operator-splitting technique. We have demonstrated that we can exclude e-/e- interactions from the N-body calculation, and then recover the physical effect by using electrostatic PIC (with the ions excluded). This changes the run-time scaling from $(N_e + N_{ion})^2$ to $N_e * N_{ion}$, which is orders-of-magnitude faster, because $N_{ion} \ll N_e$.

VORPAL shows good strong scaling behavior on Franklin to about 4,000 cores for electrostatic PIC simulations of electron cooling.

BeamBeam3D shares a common origin with IMPACT, MaryLie-IMPACT and IMPACT-T. These codes provide a sophisticated 3-D electrostatic PIC (particle-in-cell) solver with

[6] G.I. Bell *et al.*, “Simulating the dynamical friction force on ions due to a briefly co-propagating electron beam”, *J. Comput. Phys.* **227**, 87148735 (2008).

a wide range of beam dynamics and modeling capabilities. IMPACT-T (Integrated Map and Particle Accelerator Tracking-Time) is a parallel, three-dimensional, quasi-static beam dynamics code used to study dynamics in photoinjectors and RF linear accelerators. Developed under the SciDAC program, it uses a parallel, relativistic PIC method to model the space-charge forces of the beam(s) self-consistently with time as the independent variable. (The design of RF accelerators is normally performed using position, either the arc length or z-direction, rather than the time, as the independent variable.) IMPACT-T is unique in the accelerator modeling field in several ways. It solves the electro-static Poisson equation in the moving beam frame and allows an arbitrary overlap of the fields for a variety of beam-line element types. It also includes an integrated Green function method, multiple energy bins for beams with large energy spread, and models for treating RF standing wave and traveling wave structures so it can model both standing and traveling wave structure, and can include a wide variety of external magnetic focusing elements such as solenoids, dipoles, quadrupoles, etc. The integrated Green function space-charge solvers can efficiently and accurately treat beams with large aspect ratios, and a shifted Green function efficiently treats image charge effects of a cathode. It also has a flexible data structure (implemented in object-oriented Fortran) that allows particles to be stored in containers with common characteristics; for photoinjector simulations the containers represent multiple slices, but in other applications they could correspond, e.g., to particles of different species.

BeamBeam3D shares much of the software infrastructure and algorithms described above, but it has been specially developed for beam-beam simulations. It is being used to support of the electron lens experiment at BNL and also to support MEIC design studies at JLAB.

The Framework of Unified Accelerator Libraries (UAL) is a particle accelerator modeling framework that has been developed and used at BNL for over 10 years. Recently, the spin tracking code SPINK was integrated into UAL and ported to run on NVIDIA Fermi GPU hardware. This new spin-orbit tracking software, referred to as UNL-SPINK, can track 10,000 particles 100 times faster on a single 'Fermi' NVIDIA card than on a single 2.4 GHz Intel core. This enables tracking the spin and orbit for 10K particles through a typical RHIC acceleration ramp, which previously wasn't possible. UAL-SPINK is a success story for the porting of accelerator codes to the GPU, which is important in preparing for the next-generation of supercomputers.

12.2.1.3 HPC Requirements

Electron Cooling:

Simulating the modulator section of a coherent electron cooling system requires the use of delta-f PIC to control noise and the use of more than 100 macroparticles per cell to accurately capture the effects of finite temperature. Given a 3-D rectangular domain of length ~ 20 Debye units across, one obtains a mesh of $\sim 10^8$ cells with $> 10^{10}$ particles. Simulating one-half of a plasma period requires $\sim 1,000$ time steps, which implies $\sim 4,000$ processor hours for VORPAL, which are typically run on 4,000 cores. Parameter scans in

this regime over the one year period will use ~250,000 processor hours. The output from VORPAL is coupled into the FEL code GENESIS and then coupled back into VORPAL to simulate the kicker. Another 250,000 processor hours is requested for the kicker simulations. Total request for this effort is 0.5 million hours.

The large number of simulations will enable consideration of a wide array of non-ideal effects (e.g. external magnetic fields and non-Gaussian electron velocities). Because there is no theory with which to compare in such non-ideal regimes, the new 2D2V Vlasov-Poisson algorithm in VORPAL will be used to benchmark the delta-f PIC simulations. These Vlasov-Poisson runs will require a 4-D mesh, while the PIC runs will only need a 2-D mesh. Hence the runtime will be dominated by the Vlasov-Poisson runs. Based on runtimes for 2-D and 3-D charged fluid simulations with VORPAL, we estimate that a million processor hours will be required.

Beam-beam simulations:

Self-consistent beam-beam simulations are computationally demanding, requiring about 10,000 MPP hours for each run, which is in fact a parameter scan done concurrently and consists of many individual runs. A million core-hours per year would be required for parallel nonlinear optimization and system design.

Spin tracking:

UAL-SPINK is run on Carver, Magellan and Dirac machines at NERSC. Many more studies are necessary to understand RHIC operations and key sources of polarization loss. These studies would involve 8-128 cores on Carver or Magellan. It is expected that 200,000 core hours per year will be required in the near term.

12.2.1.4 Computational and Storage Requirements Summary

	Current (NERSC 2010)	2014 Requirement
Computational Hours	0.7 M	5 M
Parallel Concurrency	4k	80k
Wall Hours per Run	2	12
Aggregate Memory	5 GB	100 GB
Memory per Core	1 MB	1 MB
I/O per Run Needed	100 GB	1 TB
On-Line Storage Needed	1 TB	10 TB
Data Transfer Needed	negligible	negligible

12.2.1.5 Support Services and Software

These efforts require the following software: parallel I/O via HDF5; Trilinos; python; VisIt; and IDL. We may need assistance with visualizing 4-D and 6-D fields. Although I/O is not now a bottleneck, it does not appear to scale well and so assistance with that

may be necessary. We also know that dynamic load balancing may be required for good efficiency and that we must move to smaller surface-to-volume ratios or MPI communication-related overhead will become a bottleneck. In the future, as we move towards concurrent use of 10^5 cores, fault tolerance will become a major concern.

12.2.1.6 Emerging HPC Architectures and Programming Models

The parallel VORPAL framework has been ported run in production mode on multiple GPUs for simulating electromagnetics with complex conducting boundaries (i.e., SRF cavities). More directly relevant to efforts described above, the VORPAL development team has prototyped electrostatic PIC on the NVIDIA Fermi architecture. It is expected that production VORPAL simulations of electron cooling and related electrostatic PIC will have been ported to the GPU within the next 12 months.

As noted above, UAL-SPINK has been successfully modified to obtain 100x speedup on a single NVIDIA ‘Fermi’ GPU. Since UAL already has been written to run in parallel, there are plans to develop the capability to use multi-GPU machines to accelerate the analysis of previous and future RHIC runs via simulations of polarized tracking through the acceleration ramp.

We will also address simulation of large-scale electromagnetic structures (relevant to SRF cavities for DOE/NP applications), which is computationally very demanding. We will continue to develop GPU-based LU decomposition methods for solving large-scale method-of-moments scattering problems. One goal of the project is to determine how these methods scale with increasing number of GPUs. The double-precision capability of the new Fermi GPUs available on Dirac will allow us to optimize our codes on new hardware.

Appendix A. Attendee Biographies

David L. Bruhwiler is the Vice President for Accelerator Technology at Tech-X Corp. He has 20 years of experience in the development of algorithms and high-performance software for the design and simulation of particle accelerators and other beam and plasma devices. From 1992 through 1997, Dr. Bruhwiler designed RF photocathode electron guns and subsequent beamlines for the generation of high charge (>1 nC) sub-picosecond electron pulses. Upon joining Tech-X Corp., Dr. Bruhwiler co-developed a unique algorithm for modeling particle trajectories far from the accelerator axis. He also co-developed the electrostatic and electromagnetic PIC codes OOPIC Pro and VORPAL, including the use of MPI messaging and the parallel sparse matrix solver Aztec, as well as the co-development of multiscale algorithms. Dr. Bruhwiler has used these codes extensively to study plasma-based particle accelerator concepts. He is a member of the VORPAL development team and recent work has included implementation of a 4th-order electromagnetic update, participation in implementation of the Dey-Mittra algorithm for cut-cell boundaries and modeling photocathode electron sources. A recent focus has been the development, implementation and use of algorithms in VORPAL for modeling electron cooling physics.

Joe Carlson is Group Leader for the Nuclear and Particle Physics, Astrophysics and Cosmology Group (T-2) at the Los Alamos National Laboratory. His research centers on studies of strongly interacting quantum systems. He holds a Ph. D. and a M.S. in Physics, both from the University of Illinois at Urbana-Champaign. He was a J. Robert Oppenheimer Postdoctoral Fellow at Los Alamos from 1986-1989 and became a member of the technical staff at Los Alamos in 1989. He is also an American Physical Society Fellow.

Robert Edwards received a B.S in Physics and a B.S. in Mathematics in 1984 from the Univ. of Texas at Austin. He obtained his Ph.D. in Physics in 1989 from New York University. He was a postdoc and later a Staff Scientist at the Supercomputer Computations Research Institute (SCRI), Florida State Univ., from 1989 to 1999, where he was a Co-PI on the Theoretical High Energy Physics grant (DOE). Research included development of new algorithms for calculations within spin systems, and for dynamical fermion calculations within lattice QCD as part of the HEMCGC grand challenge project. Also, new methods were developed for calculations with chiral fermion actions which elucidated the role of topology within QCD. In 1998, he shared the Gordon Bell Prize for the development of the QCDSF supercomputer. Since 1999, he has been a Staff Scientist in the Theory Group at Jefferson Lab. During this time, he jointly started the lattice group within JLab. He has led the lab's effort in developing the infrastructure for the DOE SciDAC program, in particular, for the USQCD collaboration of which he a member of the Software Committee and the Program Allocations Committee. Current research involves the determination of the highly excited state spectrum of hadrons within QCD as part of the Hadron Spectroscopy Collaboration. This work is using the DOE INCITE computing facilities at ORNL and ANL, as well as NSF, LANL and NERSC facilities.

Jonathan Engel is Professor and Associate Chair of the Department of Physics and Astronomy at the University of North Carolina, Chapel Hill. His research is in nuclear theory and its application to problems of fundamental importance in particle physics and astrophysics. Topics in this interdisciplinary line of research include double beta decay (and the mass of the neutrino), Parity and CP violation in nuclei, nucleosynthesis in stars and supernovae, neutrino scattering and the weak interaction in nuclei. He is also interested in the foundations of nuclear structure, particularly density-functional theory in heavy nuclei.

Richard Jones is Associate Professor of Physics at the University of Connecticut. His research interests are in hadron spectroscopy, strong interaction phenomenology, and parity-violating electron scattering. He holds a Ph.D. in Physics from Virginia Polytechnic Institute.

Daniel Kasen is Assistant Professor of Astronomy and Assistant Professor of Physics at UC Berkeley. Prof. Kasen's interests are in theoretical and computational astrophysics, with an emphasis on supernovae and other energetic transients. He uses multi-dimensional supercomputer simulations to study astronomical explosions and their applicability as probes of cosmology and fundamental physics. More generally, he is interested in radiation transport across a range of astrophysical environments, from galaxies to extrasolar planets. Prof. Kasen received his B.S. from Stanford University and his M.S. and Ph.D. in physics from UC Berkeley. Prior to returning to Cal, he was the Alan C. Davis fellow at Johns Hopkins University and a Hubble fellow at UC Santa Cruz. He joined the Berkeley astronomy faculty in 2010, jointly appointed with the physics department and the nuclear science division at LBNL.

Esmond Ng is a Senior Scientist at Lawrence Berkeley National Laboratory and head of the Applied Mathematics and Scientific Computing Department in LBNL's Computational Research Division. Prior to joining LBNL, he was a Senior Staff Member at Oak Ridge National Laboratory. Esmond's research interests include sparse matrix computation, numerical linear algebra, parallel computing, computational complexity, and mathematical software. He has been a key contributor in the DOE SciDAC program and working closely with application scientists. Esmond received his Ph.D. in Computer Science from the University of Waterloo.

Tomasz Plewa is Associate Professor of Computational Astrophysics in the Department of Scientific Computing, Florida State University. He obtained his Ph.D. degree in theoretical astrophysics from the Warsaw University and spent several years at the Max-Planck-Institut fuer Astrophysik, Garching, where he worked on the development of adaptive methods in application to astrophysical flows. His research focuses on development of theoretical and computational models of core-collapse and thermonuclear supernovae, mixing processes, high-energy density laboratory astrophysics, code validation, and development of high-performance multiphysics simulation tools. His recent focus is on design of scaled, high-energy density supernova experiments at the National Ignition Facility.

R. Jefferson (“Jeff”) Porter currently works to support computing needs of experimental High Energy Nuclear Physics (HENP) groups at NERSC and is leading a project to build an ALICE production grid facility in the US, co-located at NERSC/PDSF and Lawrence Livermore National Laboratory. Jeff earned his PhD in Nuclear Physics from UC Davis and has been a member of several HENP experiments: DLS at LBNL, NA49 and ALICE at CERN and STAR at BNL's Relativistic Heavy Ion Collider (RHIC). As the first liaison to NERSC for the Nuclear Science Division (NSD), Jeff supported initial operation of the NERSC PDSF cluster and participated in the HENP Grand Challenge Collaboration, developing data mining tools to access large data sets from the RHIC experiments. Jeff spent several years as Database and Deputy Computing Leader for the STAR experiment at BNL before returning to the west coast with the STAR group at the University of Washington. Jeff rejoined LBNL in 2006 to work with the Open Science Grid (OSG), operating testbeds for OSG and supporting deployment of grid middleware on NERSC systems. In 2009, he began to directly support ALICE and STAR computing operations at NERSC for the NSD.

Martin Savage is a Professor of Physics in the Nuclear Theory Group at the University of Washington. Martin's current research interests center around using the numerical technique of Lattice QCD to calculate the properties and interactions of nuclei directly from quantum chromodynamics. Martin earned his BSc (1983) and MSc (1984) at the University of Auckland in New Zealand, and his PhD from Caltech (1990). After Postdoctoral positions at Rutgers University (1990-1991) and UC San Diego (1991-1993) he became an Assistant Professor at Carnegie Mellon University (1993-1996), and then moved to the University of Washington in 1996. He has held a SSC Fellowship and a DOE Outstanding Junior Investigator award. Martin has published papers in a number of areas of subatomic physics, including experimental nuclear physics, theoretical particle physics, theoretical nuclear physics and Lattice QCD, and is a founding member of the NPLQCD (Nuclear Physics with Lattice QCD) collaboration.

Alejandro Sonzogni is a member of the staff at the National Nuclear Data Center at Brookhaven National Laboratory, where his responsibilities include NuDat and ENSDF evaluation. He holds a Ph.D. in Physics from the University of Washington.

W.I. (Chip) Watson is manager of High-Performance Computing group in the IT Division at Jefferson Laboratory. He is an active member of the SciDAC Particle Physics Data Grid project, working with others to develop and deploy advanced data grid software. He also is the Jefferson Lab principal investigator for QCDOC with overall responsibility for all of the software and cluster work done there, and is also a member of the QCD Executive Committee.

Appendix B. Workshop Agenda

Thursday, May 26		
Time	Topic	Presenter
8:00am	Arrive, informal discussions	
8:30	Welcome, introductions, workshop goals, charge to committee	Yukiko Sekine
8:50	Workshop outline, logistics, format, procedures	Harvey Wasserman,
9:00	NP Program Office Research Directions	Ted Barnes
9:15	NERSC Role in Nuclear Physics Research	Kathy Yelick
10:10	Break	
10:30	Case Study: Lattice QCD	Robert Edwards / Martin Savage / Chip Watson
11:45	Break (Lunch available)	
12:45	Case Study: Low Energy Nuclear Physics	Joseph Carlson / Esmond Ng / Jon Engel
1:30	Case Study: Astrophysics	Dan Kasen / Tomasz Plewa
2:15	Case Study: Heavy Ion Experiments	Jeff Porter
3:00	Break	
3:30	Case Study: Medium Energy Experiments	Richard Jones
4:15	Case Study: Nuclear Data	Alejandro Sonzogni
4:45	Case Study: Accelerators	David Bruhwiler
5:15	General Discussions	
5:30	Adjourn for the day	


Friday, May 27		
8:00am	Arrive, informal discussions	
8:30	NERSC Initial Summary	Richard Gerber
9:30	Case study format review; sample case study	Harvey Wasserman
10:00	Break	
10:45	Report schedule and process	Richard Gerber
11:00	Q&A, general discussions, breakout sessions, and lunch	
1:00pm	Adjourn	

Appendix C. Abbreviations and Acronyms


ALCF	Argonne Leadership Computing Facility
AMR	Adaptive Mesh Refinement
API	Application Programming Interface
ASCR	Advanced Scientific Computing Research
AY	Allocation Year
CUDA	Compute Unified Device Architecture
EIC	Electron Ion Collider
ESnet	DOE's Energy Sciences Network
FDTD	Finite Difference Time Domain
FEL	Free Electron Laser
FEM	Finite Element Modeling
FFT	Fast Fourier Transform
FNAL	FermiLab National Accelerator Laboratory
FRIB	Facility for Rare Isotope Beams
GCR	Generalized Collisional-Radiative
GPGPU	General Purpose Graphical Processing Unit
GPU	Graphical Processing Unit
HDF	Hierarchical Data Format
HEDP	High Energy Density Physics
HPC	high-performance computing
HPSS	High Performance Storage System
I/O	input output
IDL	Interactive Data Language visualization software
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LHC	Large Hadron Collider
LLNL	Lawrence Livermore National Laboratory
MPI	Message Passing Interface
NERSC	National Energy Research Scientific Computing Center
NetCDF	Network Common Data Format
NGF	NERSC Global Filesystem
OLCF	Oak Ridge Leadership Computing Facility
ORNL	Oak Ridge National Laboratory
OS	operating system
PDE	Partial Differential Equation
PDSF	NERSC's Parallel Distributed Systems Facility
PIC	Particle In Cell
RHIC	Relativistic Heavy Ion Collider
SC	DOE's Office of Science
SciDAC	Scientific Discovery through Advanced Computing
SLAC	SLAC National Accelerator Laboratory
SN	supernova

SRF Superconducting Radio Frequency


Appendix D. About the Cover




Visualization of a sparse matrix resulting from an MFDn computation showing nonzero elements in red, potentially nonzero blocks in green, and zero blocks in blue. This is before application of a new multilevel blocking algorithm. From Philip Sternberg, et al., http://unedf.org/content/annual_mtg.php.




Visualization showing RF waves propagating through input and output ports of an accelerator structure. Also shown is the rendering of electric field lines. Image courtesy of Prof. Kwan-Liu Ma, U.C. Davis. From “Advanced Visualization Technology for Terascale Particle Accelerator Simulations,” by Kwan-Liu Ma, Greg Schussman, Brett Wilson, Kwok Ko, Ji Qiang, and Robert Ryne, Proceedings of SC2002.



High-resolution hydrodynamic model of the merging white dwarf stellar binary. Shown is a contour plot of temperature distribution in a plane of the binary. White dwarf mergers are believed to constitute an alternative formation channel of thermonuclear (Type Ia) supernovae. Image courtesy of Tomasz Plewa, Florida State University.



Graph theoretic visualization of a sparse matrix from lattice QCD using a multilevel force-directed algorithm. Image courtesy of Timothy A. Davis (University of Florida) and Yifan Hu (AT&T Labs Research). See ACM Transactions on Mathematical Software, Vol. 38, No. 1, 2011, and <http://www.cise.ufl.edu/research/sparse/matrices/>.



One of the first full-energy collisions between gold ions at Brookhaven Lab's Relativistic Heavy Ion Collider, as captured by the Solenoidal Tracker At RHIC (STAR) detector. The tracks indicate the paths taken by thousands of subatomic particles produced in the collisions as they pass through the STAR Time Projection Chamber, a large, 3-D digital camera.

