





NP

NUCLEAR PHYSICS




EXASCALE REQUIREMENTS REVIEW



An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and Nuclear Physics

JUNE 15-17, 2016

**GAITHERSBURG,
MARYLAND**



On the cover:

Top left: Quantum fluctuations of gluons captured in a field configuration. The solid regions indicate enhanced action density, and the vectors show the space-time orientation of one of the gluon fields.

Top center: The 750-kg tellurium dioxide (TeO_2) bolometer array assembled by the CUORE Collaboration to search for the $0\nu\beta\beta$ -decay mode of nuclei.

Top right: The collision of two neutron stars (1.3 and 1.4 solar masses). Shown is a volume rendering of the temperature at $t = 8.32$ ms after simulation start.

Bottom: A cartoon illustrating that the global properties of nuclei can be determined through numerical simulation.

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.

NUCLEAR PHYSICS

U.S. DEPARTMENT OF ENERGY

An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and Nuclear Physics

Meeting Co-Chairs

Joseph Carlson, Los Alamos National Laboratory
Martin J. Savage, Institute for Nuclear Theory, University of Washington

Meeting Organizers

Richard Gerber, National Energy Research Scientific Computing Center
James Osborn, Argonne Leadership Computing Facility
Katherine Riley, Argonne Leadership Computing Facility
Tjerk Straatsma, Oak Ridge Leadership Computing Facility
Jack Wells, Oak Ridge Leadership Computing Facility

Report Section Leads

Nuclear Astrophysics

George Fuller, University of California, San Diego
Bronson Messer, Oak Ridge National Laboratory

Experiment and Data

Amber Boehnlein, Thomas Jefferson National Accelerator Facility
Jason Detwiler, University of Washington
Paul Mantica, Michigan State University
Jeff Porter, Lawrence Berkeley National Laboratory

Nuclear Structure and Reactions

Joseph Carlson, Los Alamos National Laboratory
Witek Nazarewicz, Michigan State University

Cold Quantum Chromodynamics

Robert Edwards, Thomas Jefferson National Accelerator Facility
Martin J. Savage, Institute for Nuclear Theory, University of Washington

Hot Quantum Chromodynamics

Peter Petreczky, Brookhaven National Laboratory

Michael Strickland, Kent State University

Path Forward

Katie Antypas, National Energy Research Scientific Computing Center

Richard Gerber, National Energy Research Scientific Computing Center

Katherine Riley, Argonne Leadership Computing Facility

Tjerk Straatsma, Oak Ridge Leadership Computing Facility

Sponsors and Representatives

Ted Barnes, DOE/Office of Nuclear Physics

Carolyn Lauzon, DOE/Office of Advanced Scientific Computing Research

Additional Contributors

Argonne National Laboratory Document Design and Development Team

Mary Fitzpatrick, Andrea Manning, Michele Nelson, and Vicki Skonicki, Communications, Education, and Public Affairs; Beth Cerny and Richard Coffey, Argonne Leadership Computing Facility

CONTRIBUTORS (WHITE PAPERS AND CASE STUDIES)

Harut Avakian, Thomas Jefferson National Accelerator Facility²
Yassid Ayyad, National Superconducting Cyclotron Laboratory, Michigan State University²
Steffen A. Bass, Duke University²
Daniel Bazin, National Superconducting Cyclotron Laboratory, Michigan State University²
Amber Boehnlein, Thomas Jefferson National Accelerator Facility¹
Georg Bollen, Facility for Rare Isotope Beams, Michigan State University²
Leah J. Broussard, Oak Ridge National Laboratory²
Alan Calder, Stony Brook University²
Joseph Carlson, Los Alamos National Laboratory³
Sean Couch, Michigan State University²
Aaron Couture, Los Alamos National Laboratory²
Mario Cromaz, Lawrence Berkeley National Laboratory²
William Detmold, Massachusetts Institute of Technology²
Jason Detwiler, University of Washington¹
Huaiyu Duan, University of New Mexico²
Robert Edwards, Thomas Jefferson National Accelerator Facility³
Jonathan Engel, University of North Carolina, Chapel Hill²
Chris Fryer, Los Alamos National Laboratory²
George M. Fuller, University of California-San Diego²
Stefano Gandolfi, Los Alamos National Laboratory²
Gagik Gavalian, Thomas Jefferson National Accelerator Facility²
Dali Georgobiani, Facility for Rare Isotope Beams, Michigan State University²
Rajan Gupta, Los Alamos National Laboratory²
Vardan Gyurjyan, Thomas Jefferson National Accelerator Facility²
Marc Hausmann, Facility for Rare Isotope Beams, Michigan State University²
Graham Heyes, Thomas Jefferson National Accelerator Facility¹
W. Raph Hix, Oak Ridge National Laboratory²
Mark Ito, Thomas Jefferson National Accelerator Facility²
Gustav Jansen, Oak Ridge National Laboratory²
Richard Jones, University of Connecticut¹
Balint Joo, Thomas Jefferson National Accelerator Facility²
Jérôme Lauret, Brookhaven National Laboratory²
David Lawrence, Thomas Jefferson National Accelerator Facility²
Huey-Wen Lin, Michigan State University²
Meifeng Lin, Brookhaven National Laboratory²
Olaf Kaczmarek, Bielefeld University (Germany)²
Dan Kasen, Lawrence Berkeley National Laboratory²
Mikhail Kostin, Facility for Rare Isotope Beams, Michigan State University²
Thorsten Kurth, National Energy Research Scientific Computing Center
and Lawrence Berkeley National Laboratory²
Peter Maris, Iowa State University²
Paul Mantica, Michigan State University¹
Bronson Messer, Oak Ridge National Laboratory²
Wolfgang Mittig, National Superconducting Cyclotron Laboratory, Michigan State University²
Shea Mosby, Los Alamos National Laboratory²
Swagato Mukherjee, Brookhaven National Laboratory²
Hai Ah Nam, Los Alamos National Laboratory²
Petr Navrátil, Tri-University Meson Facility, University of British Columbia²
Witek Nazarewicz, Michigan State University³

Esmond Ng, Lawrence Berkeley National Laboratory²
Tommy O'Donnell, Virginia Polytechnic Institute and State University²
Konstantinos Orginos, College of William and Mary/
Thomas Jefferson National Accelerator Facility²
Frederique Pellemoine, Facility for Rare Isotope Beams, Michigan State University²
Peter Petreczky, Brookhaven National Laboratory¹
Steven C. Pieper, Argonne National Laboratory²
Christopher H. Pinkenburg, Brookhaven National Laboratory²
Brad Plaster, University of Kentucky²
R. Jefferson Porter, Lawrence Berkeley National Laboratory³
Mauricio Portillo, Facility for Rare Isotope Beams, Michigan State University²
Scott Pratt, Michigan State University²
Martin L. Purschke, Brookhaven National Laboratory²
Ji Qiang, Lawrence Berkeley National Laboratory¹
Sofia Quaglioni, Lawrence Livermore National Laboratory²
David Richards, Thomas Jefferson National Accelerator Facility²
Yves Roblin, Thomas Jefferson National Accelerator Facility²
Martin Savage, Institute for Nuclear Theory, University of Washington³
Björn Schenke, Brookhaven National Laboratory²
Rocco Schiavilla, Old Dominion University, Thomas Jefferson National Accelerator Facility²
Sören Schlichting, Brookhaven National Laboratory²
Nicolas Schunck, Lawrence Livermore National Laboratory²
Patrick Steinbrecher, Brookhaven National Laboratory²
Michael Strickland, Kent State University¹
Sergey Syritsyn, Stony Brook University²
Balsa Terzic, Old Dominion University²
Robert Varner, Oak Ridge National Laboratory²
James Vary, Iowa State University²
Stefan Wild, Argonne National Laboratory²
Frank Winter, Thomas Jefferson National Accelerator Facility²
Remco Zegers, National Superconducting Cyclotron Laboratory, Michigan State University²
He Zhang, Thomas Jefferson National Accelerator Facility²
Veronique Ziegler, Thomas Jefferson National Accelerator Facility²
Michael Zingale, Stony Brook University²

¹Contributed to white paper.

²Contributed to case study.

³Contributed to both white paper and case study.

TABLE OF CONTENTS

Executive Summary	1
Abstract	1
ES.1 Summary and Key Findings	2
ES.2 Nuclear Physics Vision and Grand Challenges	3
ES.3 Priority Research Directions	4
ES.3.1 Nuclear Astrophysics	4
ES.3.2 Experiment and Data	5
ES.3.3 Nuclear Structure and Reactions	5
ES.3.4 Cold Quantum Chromodynamics	7
ES.3.5 Hot Quantum Chromodynamics	8
ES.4 Path Forward	9
1 Introduction	11
1.1 Goal of the DOE Exascale Requirements Reviews	11
1.1.1 Previous DOE Requirements-Gathering Efforts: “Lead with the Science”	11
1.1.2 National Strategic Computing Initiative	11
1.2 Nuclear Physics Exascale Requirements Review	12
1.3 Report Organization	13
2 Nuclear Physics: Vision and Grand Challenges in an Exascale Computing Ecosystem	15
2.1 Nuclear Physics Vision	16
2.2 Nuclear Physics Grand Challenges	18
2.2.1 How Did Visible Matter Come into Being and How Does It Evolve?	18
2.2.2 How Does Subatomic Matter Organize Itself and What Phenomena Emerge?	18
2.2.3 Are the Fundamental Interactions Basic to the Structure of Matter Fully Understood?	19
2.2.4 How Can the Knowledge and Technological Progress Provided by Nuclear Physics Best Be Used to Benefit Society?	20
2.3 Mapping the Nuclear Physics Grand Challenges to the Exascale Computing Ecosystem	21
2.3.1 Nuclear Astrophysics	22
2.3.2 Experiment and Data	22
2.3.3 Nuclear Structure and Reactions	23
2.3.4 Cold Quantum Chromodynamics	24
2.3.5 Hot Quantum Chromodynamics	25

2.4	ASCR, NSF, and Nuclear Physics.....	26
2.4.1	Present Day	26
2.4.2	Going Forward	28
3	Research Directions and Computing Needs/Requirements.....	31
3.1	Nuclear Astrophysics	31
3.1.1	Scientific Challenges and Opportunities	31
3.1.2	Priority Research Directions	31
3.1.3	Cross-Cutting Research Directions.....	41
3.1.4	Computing, Data, and Networking Needs and Requirements.....	42
3.2	Experiment and Data.....	45
3.2.1	Scientific Challenges and Opportunities	45
3.2.2	Priority Research Directions	47
3.2.3	Cross-Cutting Directions.....	53
3.2.4	Computing and Data Needs	56
3.3	Nuclear Structure and Reactions.....	57
3.3.1	Scientific Challenges and Opportunities	57
3.3.2	Priority Research Directions	60
3.3.3	Cross-Cutting Issues and Opportunities.....	66
3.3.4	Computing Needs and Requirements	67
3.4	Cold Quantum Chromodynamics.....	69
3.4.1	Scientific Challenges and Opportunities	69
3.4.2	Priority Research Directions	72
3.4.3	Cross-Cutting Research Directions	81
3.4.4	Computing Needs and Requirements	83
3.5	Hot Quantum Chromodynamics.....	88
3.5.1	Scientific Challenges and Opportunities	88
3.5.2	Priority Research Directions	89
3.5.3	Cross-Cutting Research Directions.....	94
3.5.4	Computing Needs and Requirements	95
4	Path Forward.....	97
4.1	Designing an Exascale Ecosystem for Nuclear Physics.....	97
4.2	Summary	101
5	References.....	103
6	Acronyms and Abbreviations.....	109

Appendices: Meeting Materials.....	113
Appendix A: Nuclear Physics Organizing Committee and Meeting Participants.....	A-1
Appendix B: Nuclear Physics Meeting Agenda.....	B-1
Appendix C: Nuclear Physics White Papers.....	C-1
Appendix D: Nuclear Physics Case Studies.....	D-1

FIGURES

2-1	Computational facilities, network connections, and nuclear physics-relevant facilities critical to the national nuclear physics research program.	16
2-2	How did visible matter come into being and how does it evolve?	18
2-3	Phases of nuclear matter in the crust of a NS.	19
2-4	Two scenarios exist for the mass ordering, or hierarchy, of the neutrinos.	20
2-5	The organization of SciDAC collaborations between NP and ASCR scientists are shown for the USQCD NP project and for the NUCLEI project.	27
2-6	The capability and capacity computing resources available to the USQCD NP SciDAC project and the INCITE awards associated with the NUCLEI SciDAC project.	27
3-1	The specific entropy of the stellar material from a 3-D CCSNe simulation by Lentz et al.	32
3-2	The collision of two NSs.	33
3-3	Coherent neutrino oscillations: transformation probabilities and final evolved spectra compared with initial spectra.	35
3-4	R-process nucleosynthesis including simplified single-angle calculations of coherent neutrino oscillations, and more complete multi-angle calculations.	36
3-5	Potential impact of neutrino-matter interactions in CCSNe – the so-called ‘halo’ effect.	37
3-6	A volume rendering of radial velocities for some representative models of helium burning on the surface of white dwarfs.	39
3-7	Vertical velocity showing the convective structure in a low-mach XRB simulation.	41

3-8	The capability and capacity computing resource requirements and the hot and cold data requirements in 2025 to accomplish the science objectives of the Nuclear Astrophysics program.	44
3-9	The 750-kilogram tellurium dioxide bolometer array assembled by the CUORE Collaboration to search for the $0\nu\beta\beta$ -decay mode of nuclei.	46
3-10	GRETINA is the predecessor for the GRETA gamma-ray tracking detector array.	49
3-11	The High-Resolution Spectrograph is one of the planned spectrometers for FRIB.	49
3-12	An ALICE event display of a real Pb-Pb collision at 5.02-tera electron volt nucleon center of mass energy.	51
3-13	JLab aerial view after the \$338M 12-GeV upgrade showing the expanded experimental facilities.	55
3-14	The capability and capacity computing resource requirements and the hot and cold data requirements in 2025 to accomplish the objectives of the Experiment and Data program.	56
3-15	Case studies of the NS&R exascale program; the basic elements that atomic nuclei are made of depend on the energy of the experimental probe and the distance scale.	58
3-16	Collaboration links within NUCLEI SciDAC, illustrating the integration of the project as of 2016.	59
3-17	Schematic illustration of the scientific method as applied to NS&R research.	60
3-18	Nuclei calculated in ab initio approaches prior to 2005 and now.	61
3-19	Capture of ^3He by alpha particles to form ^7Be	63
3-20	Success and challenges in electroweak nuclear physics.	64
3-21	Spontaneous fission yields of ^{240}Pu predicted in nuclear DFT.	65
3-22	Mass-radius relation for NSs.	66
3-23	The Quantum Ladder: physical systems at various scales, from microscopic to macroscopic.	67
3-24	The capability and capacity computing resource requirements and the hot and cold data requirements in 2025 to accomplish the science objectives of the NS&R program.	68

3-25	Cover of <i>Nature</i> magazine containing an article that discusses the conflicting experimental measurements of the proton radius.	69
3-26	Gravitational waves emerging from inspiraling of binary NSs depend upon the nuclear EoS.	70
3-27	Cover of the NSAC <i>2015 Long-Range Plan for Nuclear Physics</i>	71
3-28	Lattice QCD calculations of the gluonic structure of the proton that are anticipated in the exascale era will support and complement an EIC experimental program envisioned as the future of QCD research in the United States.	73
3-29	Present-day calculations of the proton TMDs for the up and down quarks at $x=0.1$	73
3-30	The structure of conventional and hybrid mesons.	75
3-31	Isoscalar and isovector meson spectra at a pion mass of 391 MeV and spatial volume of 3 fm.	75
3-32	The spectrum of light nuclei and hypernuclei when the up-, down-, and strange-quarks have equal mass.	77
3-33	The magnetic moments of the light nuclei at the SU(3) symmetric point and the corresponding experimental values	77
3-34	A schematic of the isotope separation at FRIB.	78
3-35	Observation of a nEDM would be evidence that our universe does not respect time-reversal symmetry.	79
3-36	The quantum fluctuations of the gluons fields captured in a field configuration in one ensemble.	80
3-37	Schematic of the current USQCD software stack.	82
3-38	The present-day linkages between ASCR researchers and domain scientists in cold QCD.	83
3-39	A schematic of the three-phase work flow currently used in cold QCD calculations for nuclear physics observables.	84
3-40	The capability and capacity computing resource requirements and the hot and cold data requirements in 2025 to accomplish the science objectives of the Cold QCD program....	85
3-41	Finite-temperature and baryochemical potential lattice QCD calculations allow the exploration of the transition and crossover lines separating the low-temperature hadronic phase from the QGP phase.	90

3-42 Nuclear suppression factor R_{AA} for the $\Upsilon(1s)$ and $\Upsilon(2s)$ bottomonia states as a function of the number of nucleons participating in the heavy ion collision and particle transverse momentum.91

3-43 Visualization of an IP-Glasma-based EBE initial condition for the energy density distribution generated in a high-energy heavy-ion collision, along with the evolved distribution at $\tau = 6$ fm/c using two values of the shear viscosity to entropy density ratio: $\eta/s = 0$ and $\eta/s = 0.16$92

3-44 The capability and capacity computing resource requirements and the hot and cold data requirements in 2025 to accomplish the science objectives of the Hot QCD program.95

4-1 The projected computational resources required by nuclear physicists.98

EXECUTIVE SUMMARY

Abstract

Imagine being able to predict — with unprecedented accuracy and precision — the structure of the proton and neutron, and the forces between them, directly from the dynamics of quarks and gluons, and then using this information in calculations of the structure and reactions of atomic nuclei and of the properties of dense neutron stars (NSs). Also imagine discovering new and exotic states of matter, and new laws of nature, by being able to collect more experimental data than we dream possible today, analyzing it in real time to feed back into an experiment, and curating the data with full tracking capabilities and with fully distributed data mining capabilities. Making this vision a reality would improve basic scientific understanding, enabling us to precisely calculate,

Imagine being able to predict — with unprecedented accuracy and precision — the structure of the proton and neutron, and the forces between them, directly from the dynamics of quarks and gluons, and then using this information in calculations of the structure and reactions of atomic nuclei and of the properties of dense neutron stars...

for example, the spectrum of gravity waves emitted during NS coalescence, and would have important societal applications in nuclear energy research, stockpile stewardship, and other areas. This review presents the components and characteristics of the exascale computing ecosystems necessary to realize this vision.

Nuclear physics research and its applications, more than many other areas of science, rely heavily on large-scale, high-performance computing (HPC). HPC is integral to (1) the design and optimization of an extensive and vibrant experimental program, (2) acquisition and handling of large volumes of experimental data, and (3) large-scale simulations of emergent complex systems — from the subatomic to the cosmological. Dramatic progress has already been made in enhancing our understanding of strongly-interacting matter across many areas of physics: from the quark and gluon structure of the proton and neutron, the building blocks of atomic nuclei; to relating the hot dense plasma in the early universe to the near-perfect fluid produced at the Relativistic Heavy Ion Collider (RHIC); to the structure of rare isotopes and the reactions required to form all the elements of the universe; to the creation and properties of the dense cold matter formed in supernovae and NSs.

In addition to pursuing new analytical and experimental techniques and algorithms and more complete physical models at higher resolution, the following broadly grouped areas relevant to the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) and the National Science Foundation (NSF) would directly affect the mission need of the DOE Office of Science (SC) Nuclear Physics (NP) program.

- Exascale capability computing and associated capacity computing will revolutionize our understanding of nuclear physics and nuclear applications.
- Closely tied to the needs for exascale hardware are the requirements for new software (codes, algorithms, and workflows) appropriate for the exascale ecosystem, which can be developed through collaborations among ASCR and NSF mathematicians and computer scientists and NP scientists.
- There is a need to read, write, manage, analyze, curate, and track data of a complexity and scale never before encountered. This need is important in all areas of nuclear physics.
- Growing and sustaining a workforce to carry nuclear physics through the exascale era is vital to meeting the NP mission; workforce development should include enhanced collaboration between NP, ASCR, and NSF and, in particular, new positions at the ASCR/NP interface.

ES.1 Summary and Key Findings

The mission of the NP program is to discover, explore, and understand all forms of nuclear strongly interacting matter. The fundamental degrees of freedom in nuclear physics are quarks and gluons, whose dynamics create diverse emergent phenomena and nuclear matter over an enormous range of length scales and energies. HPC plays an essential role in understanding these emergent systems and in establishing predictive capabilities in the associated energy regimes. In this review, NP scientists focused on the scientific challenges in the areas of Nuclear Astrophysics, Experiment and Data, Nuclear Structure and Reactions, Cold Quantum Chromodynamics (QCD), and Hot QCD. These areas pertain to understanding the high-temperature quark-gluon plasma (QGP) in the early universe; the collective QCD phenomena that form nucleons, describe their interactions, and provide laboratories to probe the fundamental laws of nature; the structure and reactions of a diverse range of atomic nuclei; the evolution of the dense matter formed in supernovae and NSs; and environments that create the elements in the universe. Scientists in each NP area have worked extensively to maximally exploit the leading computational and software capabilities to advance nuclear science. The computational, data, software, and workforce needs were evaluated in terms of their ability to meet the mission of the NP program. Many of the requirements identified in this report have their precursors in previous reports, including the original *Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale* meeting (DOE 2009), the National Academies decadal study report (National Academies Press 2013), and the 2015 *Long Range Plan for Nuclear Science* (Long-Range Plan) (DOE and NSF 2015).

Wide availability of exascale computing is key to critical progress in nuclear physics across all areas. Software, algorithm, and workflow development are required to allow NP to take advantage of exascale computing architectures.

- Nuclear scientists need sustainable and high-performance software tools, programming models, and applications to enable effective use of exascale systems.
- New algorithms are needed to enable codes to run efficiently on upcoming HPC architectures, allowing scientists to model larger systems at higher resolution with greater accuracy.
- New applications are required to perform the suites of simulations necessary to explore new parameter regimes and to quantify uncertainties.
- New hardware and software tools are needed to analyze, track, and manage data produced by experiments and in simulations, including developments in databases, and to move data efficiently between sites for appropriate analysis.


The NP community is convinced that advanced computing, mathematics, and computer science can have a transformational impact on nuclear physics.

The growing needs of NP experiments, and the data-related needs of the community in general, have been identified and characterized. The data analysis, storage, archiving, curation, movement, and tracking requirements in both the experimental and theoretical components of the NP program are increasing very rapidly. Further, the experimental program would benefit from real-time access to the advanced computing capabilities of ASCR and NSF.

An exascale ecosystem for nuclear physics requires a well-trained, diverse, and capable workforce to achieve the vision outlined in this review. Such a workforce is crucial to meeting the scientific challenges that face nuclear physics, and it should include applied mathematicians, computer scientists, statisticians, and computational nuclear physicists. Particularly critical for the long-term success of this program is the growth in the number of permanent positions for scientists working at the interface of applied math and computer science with NP, both at national laboratories

and universities. Successful and long-term career paths for such scientists are only now being developed, and further advances require a strong coordination and collaboration between ASCR, NSF, and NP.

Participants in the NP Exascale Requirements Review were asked to identify the grand challenges facing nuclear physics, determine the priority research topics and computing needs required to successfully meet these grand challenges, and outline the most important steps in the path forward to reach the goals.



Results of simulating a Type-1a supernova from merging white dwarfs, each with 1.06 solar mass. (Moll et al. 2014)

ES.2 Nuclear Physics Vision and Grand Challenges

Nuclear physics is a diverse field with strong internal connections and with external relations to many areas of science — from the hot dense matter of the early universe and the properties of protons at the smallest length scales, to the properties of nuclei and the reactions that form all the elements, to the dense matter and violent explosions of NS mergers and supernovae. It reaches from the most fundamental physics, including searches for physics beyond the Standard Model and the properties of neutrinos, to critical societal problems including the production of energy through nuclear fission and fusion and applications in nuclear medicine. The overarching questions addressed by the nuclear science community have been succinctly formulated in the National Academies nuclear physics decadal study report (National Academies Press 2013):

1. How did visible matter come into being and how does it evolve?
2. How does subatomic matter organize itself and what phenomena emerge?
3. Are the fundamental interactions that are basic to the structure of matter fully understood?
4. How can the knowledge and technical progress provided by nuclear physics best be used to benefit society?

Computational nuclear physics plays a key role in answering each of these questions. It addresses, both through large-scale simulations of astrophysical objects and through studies of nuclei and their reactions, the questions of how all the elements came into being. Computational nuclear physics is required to interpret experimental data and thus return maximum results from DOE investments in experimental facilities, and to translate these results into new knowledge about subjects as diverse as the quark-gluon liquid created at the RHIC at Brookhaven National Laboratory (BNL), the quark and gluon structure of protons, the interaction between nucleons arising from these structures, and the properties of nuclei and nucleonic matter arising from the interaction between nucleons.

For all these reasons, computational nuclear physics was highlighted in all areas in the community process leading up to the 2015 Long-Range Plan, which states: “We recommend new investments in computational nuclear theory that exploit the U.S. leadership in high-performance computing” (DOE and NSF 2015).

Programs such as Scientific Discovery through Advanced Computing (SciDAC) and the Exascale Computing Project (ECP) are essential to these new transformational opportunities in nuclear science. Revolutionary advances in computing, mathematics, algorithms, and data will dramatically alter the field of nuclear science. DOE SC envisions reaching the exascale level of computing within the next decade (by 2026). With simultaneous advances in applied mathematics, computer science, software and data, and NP itself, we can expect a transformation of nuclear science with dramatic impact in all areas, including experiment and applications.

With simultaneous advances in applied mathematics, computer science, software and data, and NP itself, we can expect a transformation of nuclear science with dramatic impact in all areas, including experiment and applications.

ES.3 Priority Research Directions

The NP Exascale Requirements Review participants focused on five main areas for which exascale resources are required to achieve the goals of the nuclear physics community. There are common elements in the needs in these five areas, but there are also requirements that are unique to each.

ES.3.1 Nuclear Astrophysics

Nuclear astrophysics research offers exciting new opportunities, tying together new laboratory experiments and new astrophysical observations with dramatically increased capabilities in HPC. Cutting-edge efforts in computational nuclear astrophysics are fundamentally multi-physics in nature, and large-scale numerical simulations are at the heart of this science.

Key objectives in the Long-Range Plan (DOE and NSF 2015) include exploring the origin of the elements (e.g., ascertaining the site[s] where the heavy r-process elements, such as uranium and gold, are synthesized), the physics of ultra-dense neutron-rich nuclear matter, and the nature of neutrinos and their interactions with nuclei and in dense nuclear matter.

Selected Highlights of Exascale-enabled Nuclear Astrophysics

- Complete three-dimensional calculations of NS mergers, including realistic treatments of strong interaction and neutrino microphysics and nucleosynthesis including uncertainties.
- Complete three-dimensional plus temporal evolution of core collapse supernovae (CCSNe) with realistic microphysics to determine explosion mechanism, neutrino physics, and nucleosynthesis.
- Calculate flame spreading across the NS surface in x-ray bursts (XRBs), including a robust subgrid flame model and coupling to photon transport.
- Complete more realistic studies of neutrino quantum kinetic equations in both CCSNe and NS mergers to address potential quantum impacts in astrophysical evolution and neutrino physics.

Exascale computational resources are required to meet these objectives through large-scale simulations, demanding expanded access to new architectures. Advanced software and analysis techniques are also required, because the simulations need extensive parameter studies. Unique and challenging new kinds of physics will also be modeled — for example, neutrino flavor evolution with neutrino scattering and the gravitational signals from binary NS mergers. Another example is modeling the electromagnetic signals accompanying binary NS mergers, or the r-process nucleosynthesis of CCSNe — a key part of “multi-messenger” probes of these events. We can expect exciting new developments along many of these fronts in the exascale era.

ES.3.2 Experiment and Data

Experimental nuclear physics computing and data needs were explicitly considered in this review, a new arena of intersection between ASCR and NP. Nuclear physics experiments are driven by precision to access the multi-dimensional and multi-channel problem space, requiring beam intensity, polarization, and careful treatment of backgrounds and systematics. DOE NP experimental efforts at major facilities include nuclear studies at the Argonne Tandem Linac Accelerator System (ATLAS) at Argonne National Laboratory, electron scattering at Thomas Jefferson National Accelerator Facility (JLab), and proton and heavy-ion collisions at the RHIC at BNL. Needs for increased computational resources are being driven by the Facility for Rare Isotope Beams (FRIB), currently under construction at Michigan State University (MSU); by an anticipated Electron Ion Collider (EIC); and by fundamental physics, including the Fundamental Neutron Physics Beam Line at the Spallation Neutron Source at Oak Ridge National Laboratory and neutrinoless double-beta ($0\nu\beta\beta$) decay experiments, and research at smaller university and laboratory facilities. The ever-increasing data and analysis needs of NP experiments can be addressed, in part, through collaboration with ASCR.

Selected Highlights of Exascale-enabled Experiment and Data


- Implementation of end-to-end software frameworks that employ best practices for integrating data transport, messaging, compute services, and visualization components.
- Development of new software frameworks for fine-grained parallelism for detector simulations and data analysis.
- Creation of end-to-end software frameworks that operate at different data-processing stages to efficiently use available computing resources and adapt to real-time environments.

The major experimental needs outlined in this report are data streaming and near-line validation,¹ detector simulation and data analysis on HPC resources, the combination of work flows, data management and network infrastructure, and accelerator simulation. The capability and capacity computing needs in this area are modest compared with those in other areas, but the requirements for both immediate (hot) and long-term (cold) data are significant. Also there is a significant need to promote a close association between the experimental nuclear physics facilities and HPC resources, driving both hardware requirements in areas like networking and data storage and software efforts in data validation, analysis, and visualization. The major new nuclear physics experiments could greatly benefit from an increased connection with the forthcoming exascale ecosystem.

ES.3.3 Nuclear Structure and Reactions

Nuclear structure and reaction research is progressing dramatically as a result of exciting theoretical progress in the nuclear many-body problem and simultaneous experimental developments. The challenges driving this field include the following: determining the sites and processes creating all the elements in the universe, examining general organizational principles and processes governing the physics of nuclei, exploiting nuclei to reveal the fundamental symmetries of nature, and determining the scientific and practical uses of nuclei for society in general and for related fields of science.

¹ Near-line validation is the extraction of physics-based observables from stored data on a timescale typically within a few minutes of collection to validate correct experimental operation, enable optimization, and provide input for online data analysis.



In partnership with applied mathematics and computer science, modern nuclear structure theory strives to estimate uncertainties in predictions and assess extrapolations. Shown is the landscape of bound nuclei as a function of proton and neutron number predicted in nuclear Density Functional Theory (Erler et al. 2012). Mean drip lines, where the nuclear binding ends, and their systematic uncertainties (red) were obtained by averaging the results of different theoretical models. The inset shows the posterior distributions for the parameter set of the energy density functional UNEDF1 (McDonnell et al. 2015); this is an essential element of the statistical uncertainty estimate.

The quantum many-body methods used in nuclear physics all have analogues in other fields dealing with complex systems (e.g., condensed matter and atomic and molecular physics), but the unique features of nuclear interactions have particular requirements in the present and future eras of computational science. The requirements are, in most cases, similar to those in other areas of nuclear physics: access to significant resources on the emerging architectures and data storage and analysis needs. Some algorithms — for example, the large-scale linear algebra used in configuration interaction methods — have particular needs for large memory per node as well. Also, here and in other areas, there are significant workforce needs, particularly for permanent positions at the interface of nuclear science, computer science, and mathematics.

Selected Highlights of Exascale-enabled Nuclear Structure and Reactions

- Complete ab initio calculations of neutron-rich nuclei and electromagnetic transitions, including those to be studied at FRIB and beyond, with quantified uncertainties.
- Identify global nuclear physics inputs, with theoretical uncertainties, for r-process nuclei and develop the universal energy density function for all nuclei.
- Make quantified predictions of the alpha-capture reactions that synthesize carbon, oxygen, and heavier elements in stars.
- Calculate $\beta\beta$ decay matrix elements of ^{76}Ge , ^{136}Xe , and ^{130}Te with uncertainties of less than 25% and determine ab initio electroweak responses of ^{40}Ar and nucleonic matter in astrophysical environments.

The questions being addressed in this field are very closely tied to large-scale experiments and facilities in nuclear science and beyond. In some cases, these connections are obvious, like those in the lower-energy nuclear facilities such as ATLAS at Argonne, the National Superconducting Cyclotron Laboratory (NSCL) and FRIB at MSU, and other university laboratories. The science of nuclei also benefits greatly from connections to JLab and important accelerator neutrino experiments like the Deep Underground Neutrino Experiment (DUNE). Another major new

initiative of nuclear science is the search for $0\nu\beta\beta$ decay, the observation of which would provide conclusive proof of the violation of lepton number and the Majorana nature of the neutrino.

An improved theoretical understanding is essential in optimizing the physics resulting from the greatly enhanced experimental program. Nuclear many-body theory holds great promise for dramatically advancing our understanding and exploitation of nuclei at laboratories for fundamental physics research through exascale computing.

ES.3.4 Cold Quantum Chromodynamics


Protons, neutrons, and their interactions arise from quarks and gluons and their dynamics, as dictated by QCD. A precise understanding of the proton, the neutron, and the inter-nucleon forces is critical to many aspects of subatomic science. The three-dimensional structure of the nucleon and its excitations is being explored with precision by the experimental programs at JLab and RHIC. These programs will be extended to studies of the gluonic structure of nucleons and nuclei with construction of the planned EIC. Understanding/refining the forces between the neutrons and protons is important to building a more comprehensive picture of nuclei and their reactions; such research is a focus of the FRIB facility.

Selected Highlights of Exascale-enabled Cold QCD

- Complete precision calculations of nucleon couplings and associated form factors to all quark and gluon interactions. Obtain a detailed understanding of the gluonic contributions to the mass and spin of the proton, nuclei, and exotic states.
- Complete a quantitative exploration of the strong-interaction glue, and excitations, that binds quarks into hadrons, to support and complement an EIC experimental program.
- Perform precision calculations of light nuclei and chiral nuclear forces at the physical quark masses, including electromagnetism, to support and complement a FRIB experimental program.
- Complete precision calculations of the two- and three-nucleon short-distance interactions contributing to $0\nu\beta\beta$ -decay of nuclei and of interactions contributing to a neutron electric dipole moment (nEDM) and nuclear EDMs in support of an enhanced program in Fundamental Symmetries.

The prime areas of research in cold QCD are hadron structure and spectroscopy, nuclear interactions and structure arising from QCD, and applications of QCD to fundamental symmetries in nucleons and nuclei.

Each of these areas has close connections to the NP experimental program. HPC is essential in this effort. Indeed, the drive toward precision theoretical understanding of cold QCD systems has traditionally spurred the development of HPC. The computational requirements are large at both the capability and capacity level. Capability computing is required to generate the gauge configurations in large spacetime volumes that can later be analyzed for specific observables using capacity computing. Hot and cold data requirements are important, and workforce needs are significant for simultaneously advancing on these multiple fronts.



A lattice QCD calculation of the energy density in, and alignment of, the color fields associated with a configuration of quarks inside the proton. (Image reproduced with permission from Derek Leinweber.)

The future of lattice QCD at the exascale is very exciting. Lattice QCD will allow precise calculations of the quark and gluon structure of the proton and neutron and their excitations, and enable nucleons to be used as precise laboratories for studies of fundamental symmetries through measurements such as the nEDM. Advances into multi-nucleon systems at physical parameters, from which the nuclear forces can be refined through the use of effective field theories and modern phenomenological interactions, are important. They offer the promise of directly connecting QCD to nuclei, thereby providing a rigorous underpinning for calculations throughout the periodic table of elements and of dense hadronic matter.

ES.3.5 Hot Quantum Chromodynamics

The matter in the early universe was initially very hot and dense. As it cooled to temperatures in the vicinity of 1.5×10^{12} K, it underwent a transition from a state composed of unbound quarks and gluons — the QGP — to a gas made of hadrons, including protons, neutrons, and pions. The RHIC facility at BNL and the Large Hadron Collider (LHC) at CERN (European Organization for Nuclear Research) reverse this transition, colliding nuclei at very high energies to probe the properties of the QGP. Lattice QCD calculations use first-principles methods to study the properties of the QGP at these high temperatures. The hydrodynamics of the collisions probe the transport properties of the plasma, and they have revealed that the QGP is, in fact, nearly a perfect fluid, with an extremely low shear viscosity.

Studies of QCD at high temperatures address questions related to the nature of the equation of state (EoS) and transport properties of QCD at finite temperature and density, how the early-time dynamics of the nuclear collisions approach equilibrium conditions, and the nature of the elementary excitations in the QGP. They probe the essence of matter at the extremes of temperature and density.

Selected Highlights of Exascale-enabled Hot QCD

- Complete lattice QCD calculations of thermodynamic quantities related to RHIC and LHC heavy-ion experimental programs.
- Complete fully realistic calculations of meson spectral functions and transport coefficients.
- Achieve precise extraction of QGP properties at low and high baryon density, from first principles and from the large body of available high-precision experimental data.
- Complete fully realistic classical statistical simulations of early-time dynamics and equilibration with dynamical chiral fermions.

Hot QCD efforts span an important range of equilibrium (lattice QCD) and non-equilibrium (early-time dynamics and thermalization) efforts, leading the way to a much more complete understanding of all the properties of QCD in the hot and dense regime. The computational requirements of hot QCD emphasize the need for capability computing resources from present day through the exascale era. The data requirements are also significant, particularly for lattice QCD studies aimed at locating the QCD critical point at high temperature and relatively modest baryon density.

ES.4 Path Forward

The scientific challenges facing the U.S. nuclear physics research program are unprecedented in scope and impact. These challenges require an extensive and diverse exascale high-performance computing and data ecosystem (HPCDE). The diversity of the computing requirements emerges from the need to be able to solve QCD in both the hadronic and low-energy regimes and at extremes of temperature and density, calculate the structure and reactions of atomic nuclei, perform large-scale nuclear astrophysics simulations, and acquire and analyze large data sets generated in the nation's laboratories with real-time data-processing capabilities.

The integration of this exascale ecosystem into the nuclear physics research program will provide unprecedented predictive capabilities, sparking scientific discoveries across the wide spectrum of nuclear science, including many that we cannot imagine today.

The required exascale environment includes capability computing resources to be delivered with heterogeneous architectures and commensurate hot- and cold-data storage capabilities; enhanced infrastructure for data-handling capabilities; growth of a highly skilled and sustainable workforce (developed in close collaboration with ASCR and NSF); and increased capacity computing resources at universities and laboratories. It would be advantageous for DOE NP and NSF to develop coherent plans for nuclear physics computing in the exascale era that complement and enhance those of ASCR. The integration of this exascale ecosystem into the nuclear physics research program will provide unprecedented predictive capabilities, sparking scientific discoveries across the wide spectrum of nuclear science, including many that we cannot imagine today.

1 INTRODUCTION

1.1 Goal of the DOE Exascale Requirements Reviews

During fiscal years (FYs) 2015 and 2016, the Exascale Requirements Reviews brought key computational domain scientists, U.S. Department of Energy (DOE) planners and administrators, and experts in computer science and applied mathematics together. Meetings were held for each of the DOE's six Office of Science (SC) program offices, as follows:

- The High-Energy Physics (HEP) review was held in June 2015.
- The Basic Energy Sciences (BES) review was held November 3–5, 2015.
- The Fusion Energy Sciences (FES) review was held in January 2016.
- The Biological and Environmental Science (BER) review was held in March 2016.
- The Nuclear Physics (NP) review was held in June 2016.
- The Advanced Scientific Computing Research (ASCR) review was held in September 2016.

The overarching goal was to determine the requirements for an exascale ecosystem that includes computation, data analysis, software, work flows, high-performance computing (HPC) services, and other programmatic or technological elements that may be needed to support forefront scientific research.

Each Exascale Requirements Review resulted in a report prepared by DOE for wide distribution to subject matter experts and stakeholders at DOE's ASCR facilities, including the Argonne and Oak Ridge Leadership Computing Facility centers (ALCF and OLCF) and the National Energy Research Scientific Computing Center (NERSC).

1.1.1 Previous DOE Requirements-Gathering Efforts: “Lead with the Science”

DOE has experienced definite value in implementing its previous requirements-gathering efforts (Helland 2016). Such review meetings have served to:

- Establish requirements, capabilities, and services.
- Enable scientists, programs offices, and the facilities to have the same conversation.
- Provide a solid, fact-based foundation for service and capability investments.
- Address DOE mission goals by ensuring that DOE science is effectively supported.

1.1.2 National Strategic Computing Initiative

Dovetailing with the current Exascale Computing Project is establishment of the National Strategic Computing Initiative (NSCI) by Executive Order on July 30, 2015. NSCI has the following four guiding principles:

1. The United States must deploy and apply new HPC technologies broadly for economic competitiveness and scientific discovery.
2. The United States must foster public-private collaboration, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.

3. The United States must adopt a whole-of-government approach that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.
4. The United States must develop a comprehensive technical and scientific approach to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations (Helland 2016).

NSCI's five objectives echo plans already under way in DOE's current exascale computing initiatives. In fact, DOE is among the NSCI's three lead agencies (along with the U.S. Department of Defense and the National Science Foundation [NSF]) that recognize these agencies' historical roles in pushing the frontiers of HPC and in helping to keep the United States at the forefront of this strategically important field (Helland 2016).

1.2 Nuclear Physics Exascale Requirements Review

DOE SC convened an Exascale Requirements Review for the NP program on June 15–17 in Gaithersburg, Maryland. The review brought together leading nuclear physics researchers and program managers, scientific and HPC experts from the ASCR facilities and scientific computing research areas, and DOE NP and ASCR staff (see Appendix A for the list of participants).

Participants addressed the following:

- Identify forefront scientific challenges and opportunities in nuclear physics that could benefit from exascale computing over the next decade.
- Establish the specifics of how and why new HPC capability will address issues at various nuclear physics frontiers.
- Promote the exchange of ideas among theoretical and application scientists in nuclear physics, computer scientists, and applied mathematicians to maximize the potential for use of exascale computing to advance discovery in the field of nuclear physics.

Outlines and input from white papers and case studies (Appendices C and D, respectively), authored by the participants and submitted to the NP Organizing Committee chairs, guided the discussions in general sessions and topical breakouts. Committee members and review participants collaborated at the meeting to identify the grand challenges, priority research directions, cross-cutting research directions, and computing requirements for their fields of research — communicating these requirements to the DOE SC offices and ASCR facilities. This report therefore reflects extensive and varied forms of input from many voices in the nuclear physics community regarding HPC requirements for nuclear physics initiatives.

The review afforded a rare opportunity for the nearly 100 participants to interact and learn about each other's areas of expertise, challenges faced, and the exciting opportunities made possible by the exascale computing environment.

Exascale Requirements Reports Will Meet Multiple Needs

DOE managers will use the Exascale Requirements Review reports to guide investments and budgeting, complete their strategic planning, and respond to inquiries, including specifically their efforts to:

- Articulate the case for future needs to DOE and SC management, the Office of Management and Budget, and Congress.
- Identify emerging hardware and software needs for SC, including for research.
- Develop a strategic roadmap for the facilities based on scientific needs.

NP program managers may also use the reports to inform their work. Although balancing such varied end uses can be a challenge, the reports are intended as an information tool that can be used by many stakeholders.

1.3 Report Organization

In the balance of this Exascale Requirements Review, Section 2 provides an overview of the vision and grand challenges facing the field of nuclear physics. Section 3 addresses five areas of scientific challenge and opportunity, along with the priority and cross-cutting research directions and computing needs and requirements associated with each. Section 4 outlines a path forward for successful collaboration of the nuclear physics community with DOE's ASCR facilities (i.e., the Leadership Computing Facilities [LCFs] and NERSC). References and the acronyms/abbreviations used in the report are listed in Sections 5 and 6, respectively, followed by appendices (Organizing Committee and Meeting Participants [Appendix A], Meeting Agenda [Appendix B], White Papers [Appendix C], and Case Studies [Appendix D]).

2 NUCLEAR PHYSICS: VISION AND GRAND CHALLENGES IN AN EXASCALE COMPUTING ECOSYSTEM

This section presents an overview of nuclear physics Grand Challenges, as identified in the recent National Academies report entitled *Nuclear Physics: Exploring the Heart of Matter* (National Academies Press 2013), and a discussion of how the nuclear physics community is addressing these challenges, as described in *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015). The Grand Challenges are mapped to an exascale computing ecosystem by identifying major physics research areas and the priority research directions associated with each.

Advances in HPC, applied mathematics, computer science, and statistics are transforming research in nuclear physics and its societal applications. The nuclear physics community will require exascale computing resources to advance their work in the following critical areas:

- Essential detector and accelerator simulation;
- Data acquisition, storage, retrieval, and analysis;
- Precision calculations of nuclei and their reactions, with quantified uncertainties;
- High-fidelity studies of the nuclear astrophysical objects in our universe;
- Theoretical predictions of critical quantities directly from quantum chromodynamics (QCD) or through nuclear many-body techniques; and
- Exploration of new and exotic states of hadronic matter.


DOE's NP and the Physics Division of the NSF have extensive programs in nuclear physics research — with combined annual budgets exceeding \$500 million — that rely on HPC. These programs will depend on the exascale ecosystems that DOE's ASCR and NSF's Computer and Information Science and Engineering (CISE) division will provide. Nuclear physics researchers, in collaboration with ASCR and NSF scientists, will exploit these computing environments to meet the scientific and technological Grand Challenges facing the nuclear physics community.

As shown in Figure 2-1, DOE and NSF support a broad-based research program in theoretical and experimental nuclear physics at universities, national laboratories, and centers of excellence. They operate user facilities that can deliver highly specialized beams on designer targets, as well as sophisticated non-accelerator-based experimental facilities. These facilities are employed, for example, to precisely measure the properties and behavior of known matter, to create and explore new states of matter, and to look for new laws of nature and identify limitations of current paradigms.

DOE NP and NSF provide *capacity* computing² hardware and data storage and retrieval capabilities to support experimental and theoretical programs. ASCR, through coordination with NP, provides *capability* computing resources at leadership-class computing centers — through programs such as the ASCR Leadership Computing Challenge (ALCC) and medium-scale HPC resources at NERSC — that are critical to accomplishing the objectives of nuclear physics research. DOE and NSF also support a workforce with the specialized skill sets required to make optimal use of the computational infrastructure, through programs such as the Scientific Discovery through Advanced Computing (SciDAC) (DOE 2015) program and the newly created Exascale Computing

² Problems that require all of, or a large fraction of, one of the fastest supercomputers constitute a *capability* computing requirement. In contrast, problems that can be solved on smaller machines constitute a *capacity* computing requirement.

Project (ECP) (ECP.org, undated), in addition to providing HPC specialists with domain-science expertise who are associated with the computing facilities. ASCR provides the high-bandwidth connections that enable reliable data movement between computing facilities and scientists located at universities, national laboratories, and other research institutions.



2.1 Nuclear Physics Vision

Nuclear physics research encompasses wide-ranging and fascinating phenomena — from the behavior of quarks and gluons in the early and present-day universe; to the creation, structure, and evolution of nuclei; to the extreme astrophysical environments that created the elements. Large-scale simulations and calculations — in concert with analytical theory, terrestrial experiments, and astrophysical observation — play a unique and critical role in nuclear physics, both in developing a quantitative understanding of diverse phenomena at various energy scales and in tying them together. Such numerical calculations are critical because of the nonlinear and intrinsically quantum nature of the strong interactions, the rich set of complex nuclear and astrophysical phenomena, and the related complexity of nuclear experiments and astrophysical observations. The increasing sophistication of instrumentation, combined with advances in accelerator design, results in modern-day experiments that generate large and complex data sets. These data sets must be reliably stored and readily accessed for analyses by experimentalists distributed around the globe. Accelerators remain a central tool in experimental discovery in nuclear physics, and the future holds significant promise, with the Facility for Rare Isotope Beams (FRIB) under construction at Michigan State University (MSU) and an electron-ion collider (EIC) currently in the design phase. Large-scale simulations are essential to designing and engineering these remarkably complex devices and their associated detection instrumentation.

The transfer of technology and knowledge from nuclear physics research is accelerating and impacting a broad array of areas — from medical isotopes and health screenings, industrial uses of radioactive isotopes, production of green energy, and an array of uses in national security. Computational nuclear physics can play an important role in several applied areas. For example, nuclear theory and computation can be valuable in diagnosing the plasma formed in fusion reactors, as illustrated by efforts at the National Ignition Facility (NIF). Similarly, the applied mathematics

and algorithms developed for use in lattice QCD calculations are beginning to be transferred into other communities, such as genomics. The rapid developments in genome sequencing make it possible to use HPC implementations of linear algebra, combined with large data sets, to develop predictive capabilities for individuals.

The numerical technique of lattice QCD enables first-principles calculations of the structure and dynamics of mesons and baryons, including their spectra and underlying quark and gluon structure, as studied at Thomas Jefferson National Accelerator Laboratory (JLab) and the Relativistic Heavy-Ion Collider (RHIC) at Brookhaven National Laboratory (BNL). This technique also directly connects to experiments at RHIC that address the QCD equation of state (EoS) at both zero and finite baryon density, including future experiments to seek and examine the critical point expected at finite density. Computation and data also play an increasingly critical role in analyses of both JLab and RHIC experiments. Lattice QCD is beginning to address the physics of multi-baryon systems that are critical to our understanding of atomic nuclei and to model-independent predictions of low-energy nuclear reactions and the structure of the lightest nuclei. Scientists employ quantum many-body techniques to understand and predict the structure and dynamics of atomic nuclei, including the very neutron-rich nuclei that will be produced and studied at FRIB. An enhanced understanding of nuclei and their reactions is necessary in many applications, including diagnosing the dense plasmas formed at NIF and in astrophysical environments.

The same nuclear interactions and currents that are used in the hadronic quantum many-body problem are also employed to predict the properties of dense nucleonic matter in supernovae and neutron stars (NSs). Exascale simulations are required for a more comprehensive understanding of these environments at both the microscopic and macroscopic scales. These simulations will address critical questions including the astrophysical sites for the production of the heavy elements in the universe and the EoS of cold dense QCD, as revealed through large-scale simulations and observations of NSs and their mergers. The observation and study of gravitational waves, recently discovered at the Laser Interferometer Gravitational-Wave Observatory (LIGO), that are expected to be emitted from NS mergers will revolutionize this field.

Nuclear physics is addressing questions at the forefront of our understanding of fundamental physics dictating the emergence of our universe from the Big Bang. Nucleons and nuclei are used, for example, to probe:

- The Majorana nature of neutrinos (are they their own anti-particles?), as addressed in experimental searches for neutrinoless double-beta ($0\nu\beta\beta$) decay of nuclei — a high-priority research target in DOE NP;
- The fundamental properties of neutrinos, as measured in accelerator and reactor neutrino experiments;
- The interactions of dark matter with nuclei;
- The origin of the excess of matter over anti-matter in the universe; and
- Possible new interactions, as revealed in low-energy experiments involving neutrons and nuclei.


The community of nuclear scientists understands the ever-increasing importance of large-scale computation in addressing all these questions. *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) recognized the rising influence of these efforts and strongly recommended an increase in support for computational nuclear physics. The rapidly advancing state of the art in computational resources, math, and computer science in the exascale era will enable dramatic advances in nuclear science. An exascale computational ecosystem offers prospects for breakthroughs in all the fundamental questions the nuclear physics community is addressing — computational nuclear physics directly impacts and unifies all areas of nuclear physics.

2.2 Nuclear Physics Grand Challenges

In 2012, at the request of the National Research Council (NRC), an international committee of experts in the field of nuclear physics and other disciplines was assembled to articulate the rationale and objectives of nuclear physics research in the United States and to identify strategic future directions for the field. The resulting vision for the field is described in The National Academies report entitled *Nuclear Physics: Exploring the Heart of Matter* (National Academies Press 2013). This report emerged from the foundations established by the nuclear physics long-range planning activities that are undertaken every 5–7 years. The National Academies report identified four Grand Challenges. The following sections describe each of these challenges, followed by a discussion of the computational requirements to address each.

2.2.1 How Did Visible Matter Come into Being and How Does It Evolve?

Our universe is a mix of visible matter, dark matter, and dark energy. While we can “see” visible matter, dark energy and dark matter make themselves known only through their gravitational interactions. Much is known about the structure and properties of visible matter, but gaps remain in our understanding; such gaps limit our ability to perform calculations of subatomic quantities that directly impact society. Atoms, through complex electronic interactions, form the molecules that are responsible for life and all of the materials that drive our technologically advanced society. A complete picture of atoms requires a precise understanding of the dense, compact, and charged nuclei at their core. Because nuclei are self-bound systems of protons and neutrons, which are themselves composed of quarks and gluons, our principle challenges are understanding how protons and neutrons emerged in excess from the earliest moments of the Big Bang to form the light elements and how they are transformed into the heavy elements in stars and explosive astrophysical environments. Figure 2-2 is an illustration of the production of the elements from the earliest moments of the universe to present day. Nuclear physicists and astrophysicists are working together to address these challenges.



2.2.2 How Does Subatomic Matter Organize Itself and What Phenomena Emerge?

Quarks and gluons are not found in isolation, but instead combine themselves in quark-antiquark and three-quark clusters held together by gluons. The three-quark clusters include the proton and neutron, which are building blocks of atomic nuclei. The interactions between quarks and gluons generate most of the mass of the nucleon (proton or neutron) and also the nuclear forces between nucleons. More than 3,100 different isotopes of nuclei have been identified, and many more are anticipated to be involved in the production of elements in extreme astrophysical environments. Understanding the structure, lifetime, and reactions of these nuclei is the main objective for nuclear physics.

The complex interactions between quarks, gluons, protons, neutrons, and nuclei in general give rise to many distinct phases of matter over a range of temperatures and densities. In the earliest moments of our universe, the temperatures and densities were so high that quarks and gluons were

not confined into hadrons and, as the universe rapidly cooled, this quark-gluon plasma (QGP) transitioned through a near-perfect liquid phase, discovered at BNL's RHIC, and then into a hadronic phase prior to the production of the elements. In another extreme environment, moving out from the interior of a NS, a number of intriguing phases of matter are anticipated. In the NS's interior, novel phases of quarks and gluons may exist, possibly involving strange quarks, which transition into a more conventional nuclear matter phase at slightly lower densities. Farther out toward the crust of the NS, scientists predict quite distinct phases of nuclei such as the “pasta phases,” as shown in Figure 2-3, analogues of which have been found in biological systems. Understanding these emergent phases of strongly interacting matter that are so important in nature is essential for refining predictive capabilities in such systems.




Figure 2-3. Phases of nuclear matter in the crust of a NS. The left panel shows the result of numerical simulation revealing “nuclear pasta,” while the right panel shows 3-D printed models of the “pasta” and “waffle” phases. (Image reproduced with permission from Chuck Horowitz; available in Caplan et al. 2015).

There is significant overlap, both conceptually and practically, with other areas of science in which emergent phenomena play an important role. How strongly correlated many-body systems give rise to new phases is a central line of investigation in many important areas: traffic flow, biological systems, quantum entanglement, quantum phase transitions in condensed matter systems, nuclear matter, and the QGP.

2.2.3 Are the Fundamental Interactions Basic to the Structure of Matter Fully Understood?

Through decades of investigation, both experimentally and theoretically, the nuclear physics community now understands the fundamental interactions that dictate the nature of matter and can reliably provide quantitative descriptions of a wide range of states of matter. However, there is compelling evidence that our current theoretical description is incomplete. The Standard Model of the strong-electroweak interactions is perhaps the most successful theory that mankind has formulated. It is a unified theory of the fundamental forces resulting from a synthesis of quantum mechanics and special relativity that describes the strong nuclear force, the weak interactions, and the electromagnetic interactions between particles. However, it does not include gravity, dark matter, or dark energy. At present, dark matter and dark energy are known to exist only through

their gravitational interactions with visible matter, a force that is well described at low energies by general relativity. There are numerous experiments searching for dark matter through a variety of techniques, providing an opportunity to discover previously unidentified fundamental forces. All known matter interacts gravitationally, and the electromagnetic interaction is responsible for binding electrons to nuclei to form atoms, molecules, and solids. Nuclei and hadrons are bound together through the strong interactions, while the weak interaction — which acts over distances that are short compared to the radius of the proton — is responsible for the radioactive decay of many unstable nuclei and for the violation of certain discrete space-time symmetries. While the Standard Model can accommodate neutrino masses that are not present in its minimal formulation, little is known about their nature; experimental investigations of the properties and interactions of neutrinos offer further opportunities to identify new physics (see Figure 2-4).




Figure 2-4. Two scenarios exist for the mass ordering, or hierarchy, of the neutrinos. The absolute mass scale is not known in either scenario. (Source: DOE and NSF 2015.)

2.2.4 How Can the Knowledge and Technological Progress Provided by Nuclear Physics Best Be Used to Benefit Society?

Nuclear science has enormous impact upon modern society and plays a crucial role in attaining energy sustainability for our planet. In the decades since the discovery of the nucleus and the subsequent formulation of the fundamental forces, nuclear science has accomplished much in the areas of national security, energy, high-performance computation, health, and space exploration.

National Security

Ensuring the integrity of the U.S. nuclear stockpile is the responsibility of the National Nuclear Security Agency (NNSA) and its Stockpile Stewardship programs. Many of the scientists in these programs were trained in DOE NP and NSF research programs. The basic research efforts in nuclear science provide members of the workforce with knowledge that is central to the role of NNSA. Nuclear burning networks are at the heart of nuclear weapons explosions, and nuclear data—and the ability to predict cross sections where reliable data are unavailable—are essential in the absence of nuclear weapons testing. Similar capabilities are important to diagnose plasma conditions at NIF. The FRIB facility will make significant contributions to the required nuclear data, and anticipated advances in HPC during the exascale era will greatly enhance our predictive capabilities, as described in this report. Optimizing the utility of the FRIB program and the low-energy nuclear theory effort to national security is a challenge to the nuclear physics community.

Energy

It is imperative for our society that clean and sustainable sources of energy are deployed at scale; this research area is already a DOE priority. Research in sustainable energy sources is a broad effort, cutting across many disciplines — including chemistry and biology — and nuclear science has long played a leading role. Nuclear fission reactors are one sustainable source of energy, but they have the downside of producing long-lived nuclear waste. However, they are currently widely used in many countries, and they power, for example, cities, naval vessels, and space exploration. The anti-neutrinos produced in reactors are also used to probe fundamental neutrino properties. Stable nuclear fusion reactors are yet to be developed, and research in this area — including HPC simulations — is ongoing.

Nuclear physics and HEP lead accelerator development efforts in the United States, and both have been responsible for many Nobel Prize-winning discoveries and for much of what we currently know about subatomic physics. The ongoing transfer of this technology to other communities, such as condensed matter physics and materials, biology, and paleontology, has led to major advances in those fields as well.

Computation

While the fundamental laws of nature that dictate the strong interactions between quarks and gluons are known, solving them to predict properties of protons, neutrons, nuclei, and nuclear matter is enormously complex. Large-scale computations, requiring extreme-scale computational resources, provide solutions and insights. High-energy and nuclear physicists have contributed to the development of HPC architectures for more than 30 years. They are responsible, in part, for IBM's Blue-Genie series of supercomputers; the deployment of large, graphics processing unit (GPU)-based supercomputers, such as the one employed at JLab in 2010; and development and optimization of the requisite software. Nuclear physicists, in collaboration with computer scientists, applied mathematicians, and hardware and software vendors, continue to embrace and develop new computational architectures and develop algorithms and codes to optimally exploit the new and evolving hardware. The challenges facing nuclear physicists in the era of exascale computing are significant and will only be met through coordinated collaborations in hardware, algorithms, and coding. Such collaboration will enable the transfer of valuable technology and understanding to the commercial market, providing more effective commodity computing.

2.3 Mapping the Nuclear Physics Grand Challenges to the Exascale Computing Ecosystem

Recent developments in theoretical and experimental nuclear physics, applied mathematics, computational algorithms, machine hardware, and workforce expertise have fundamentally changed the field of nuclear physics — to the point that, with a carefully designed and deployed exascale computing ecosystem, enormous progress can be made toward accomplishing the Grand Challenges in nuclear physics research, described in Section 2.2. One objective of utilizing an exascale ecosystem is to facilitate first-principles predictive capabilities through simulations and calculations within theoretical frameworks describing physics over a large range of length scales — from quarks and gluons, to nucleons and nuclei, to explosive astrophysical environments. Such systems will allow the complete quantification of uncertainties in many critical quantities that can only be accessed through numerical calculation. A further objective is to record, analyze, and curate the large volume of data that will be produced and acquired at nuclear experimental and computational facilities. The nature of the data and the methods of acquisition and analysis are remarkably varied, and they require a diverse and complete exascale ecosystem to optimally exploit the scientific capabilities and productivity of nuclear physics facilities.

2.3.1 Nuclear Astrophysics

Nuclear astrophysics encompasses a diverse range of phenomena in extreme astrophysical environments, including supernovae and NSs. Large-scale, high-fidelity simulations offer the only opportunity to obtain sufficient understanding of these phenomena to extract fundamental knowledge from astrophysical observations. Some of the most fundamental questions about the universe, and our role in it, are within the purview of nuclear astrophysics, as summarized in the Grand Challenge question, “How did visible matter come into being and how does it evolve?”

Key questions that can be accurately addressed with exascale resources include the astrophysical sites for the formation of heavy elements in the universe, the fundamental properties and role of neutrinos in supernovae, NS cooling and NS mergers, and the emission of gravitational waves and related multi-messenger signals of NS mergers. Such studies provide probes of nuclear and high-energy physics with astrophysical observations, including neutrino properties and the QCD EoS for cold dense matter, and develop a more complete understanding of the astrophysical environments.

The increased fidelity that will become available with an exascale ecosystem is crucial to these efforts, both in terms of better treatments of the required nuclear and neutrino microphysics and higher-resolution studies in the fully relativistic, multi-dimensional radiation hydrodynamics. The floating-point-rich architectures expected in exascale machines are a key requirement in the exascale era, because they are a driving factor in reaching higher fidelity in full three-dimensional (3-D) simulations. Along with this floating-point operations requirement, sufficient memory and bandwidth are needed in an exascale ecosystem to take maximum advantage of the floating-point operations capability.

Software requirements are also critical, including the development of software, abstractions, and tools that can be portable across exascale architectures to the greatest extent possible. Similarly, the development of a diverse and capable workforce is required to explore the new, more accurate algorithms that become possible at exascale.

Large-scale simulations are essential in confronting the knowledge gained in large-scale nuclear physics laboratory experiments — including FRIB, RHIC, neutrino experiments, and others — with our present understanding of the universe around us. Exascale ecosystems offer a timely and unique opportunity to connect laboratory results to the exciting realm of astrophysical observations.

2.3.2 Experiment and Data

Nuclear physics experiments are driven by precision measurements to access the multi-dimensional and multi-channel problem space. Such experiments require beam intensity, polarization, and exquisite control of backgrounds and systematics. Over the next decade, this science will be addressed at the current and planned experimental facilities in the national nuclear physics portfolio. The successful realization of science outcomes at these facilities critically depends on the optimization of the experimental infrastructure. HPC plays an important role in the reliable and efficient execution of experiments and the rapid and well-founded analysis of resulting data. A number of areas were identified where collaboration with ASCR researchers could make a material difference in furthering the state of the art in experimental research in the next 5–10 years:

- Data streaming from detectors,
- Complex work flows that can accommodate experimental and simulations data, and
- Advanced networking and data sharing.

In addition, as in other experimental sciences, a near-term goal is to deploy petascale-class computing at beam lines. One way to begin to prepare for this deployment is to enable the experimental community to run current workloads at ASCR facilities to gain experience and

expertise. Portable software stacks and advanced scheduling mechanisms, as well as the ability to run community tools such as Geant4 and ROOT, would facilitate this capability.

The nuclear physics experimental program is marked by diversity — diversity in data rates, size of the collaboration, running periods, and complexity of the detectors. For the experimental program as a whole, storage and input/output (I/O) are primary drivers in the push toward exascale. Another driver is the capabilities enabled by matching the parallelism in detector channels to the parallelism in compute resources, which could lead to a revolution in algorithms. There is an increasing blurring between online and offline functions, and there will be an increasing emphasis on combining data from multiple experiments for global fits.

The particle accelerators at the heart of the nuclear physics user facilities rely on computationally intensive simulations to optimize accelerator performance. These include simulations of the beam transport through the machine, including the self-field interaction within the beam itself and the beam-beam interaction from the opposite colliding beam. Design of the future EIC will offer significant computational challenges. There is a near-term need to assess competing accelerator designs to identify technical pitfalls and feasibility, as well as a longer-term need for a full-fledged design study to optimize cost and performance. Both these needs can only be met with HPC, because the self-field effects and the electron cooling that need to be considered for the EIC design are based on self-consistent particle simulations that demand exascale computing.

2.3.3 Nuclear Structure and Reactions

Atomic nuclei are at the core of all visible matter and present a fascinating and challenging quantum environment to explore an extreme range of physics. Atomic nuclei are directly connected to the underlying theory of QCD through nuclear interactions that can be determined through experiment and lattice QCD calculations. They are also the basic ingredient to interpreting laboratory experiments that address questions including the stability of neutron-rich matter and the nature of the neutrino through $\beta\beta$ -decay. Atomic nuclei and nuclear matter connect laboratory experiments to the creation and evolution of nuclei in astrophysical environments, including supernovae and NSs. Large-scale simulations of atomic nuclei in an exascale ecosystem are required to address all four Grand Challenge questions (see Section 2.2).

Exascale simulations of atomic nuclei bear significant resemblance to studies of strongly correlated matter in other fields, including condensed matter and atomic, molecular, and optical (AMO) physics. Unique algorithms and understanding are required, though, because of the unique character of nuclear interactions. These interactions “tie together” spin and space through tensor and spin-orbit interactions at a much stronger level than is typical for electronic systems. Therefore, the application codes are unique to nuclear structure and reactions, although they are similar to methods used in other fields: configuration-interaction, coupled-cluster, quantum Monte Carlo, and density functional theory (DFT). While these different methods are usually employed in different regimes of nuclear masses, and for different energy and momentum regimes, a great deal has been learned by studying the interfaces between these approaches.

A common need across all these applications is dramatically increased floating-point capabilities, in both multi-core and GPU architectures. A sustained effort is required to enable these algorithms to exploit the envisioned exascale ecosystems. All algorithms need access to increased memory and I/O capabilities, and some of the algorithms, particularly configuration interaction methods, have significant memory requirements because of the importance of the three-nucleon interaction.

Software to enhance portability across GPU and multi-core architectures, such as the development of abstractions, is very important to maximize the productivity of exascale projects. Even more critical are the development and nurturing of a capable and diverse workforce, including domain

scientists, applied mathematicians and computer scientists, and a select set of people who can span both computational science and nuclear science.

Many of the highest-priority investments in nuclear science — including, for example, FRIB and searches for the $0\nu\beta\beta$ -decay of nuclei — require a sustained effort in computational nuclear science to be optimally effective. Large-scale exascale simulations offer a unique opportunity to connect physics across a wide range of scales and environments — from laboratory studies of fundamental symmetries and neutrino physics, to experiments in nuclear structure and reactions, to studies of astrophysical environments. An important objective, which will facilitate interaction with experiment, is the organization of a theoretical database that can be used, for example, to share codes and simulation results, theoretical interactions, covariances and emulators, and wave functions.

2.3.4 Cold Quantum Chromodynamics

Precise theoretical calculations that require exascale capacity and capability HPC resources, and exascale data storage and transfer capabilities are needed to accomplish the nuclear physics objectives of cold QCD. Such calculations are essential to maximizing the productivity of the nuclear physics experimental program. Because the underlying theory of strong interactions is known to be QCD, its solution through the numerical techniques of lattice QCD provides a means to robustly predict low-energy, strong-interaction properties and interactions of systems as complex as light nuclei, with a complete quantification of uncertainties. For the first time in the history of the field, a direct connection between QCD and the observed composite particles — mesons, nucleons, and nuclei — is emerging from such calculations. Accurate predictions from QCD for the properties and interactions of hadrons and nuclei will become practical. In the exascale era, these calculations will have broad impacts within and outside of nuclear physics, significantly advancing the field toward accomplishing the Grand Challenges it faces.

Underpinning all of the cold QCD calculations is the generation of samples of the strong quantum vacuum through variants of the hybrid Monte Carlo (HMC) algorithm. The required characteristics of these gauge-field configurations demand capability exascale computing resources. These configurations will be archived after production for multiple subsequent usages, requiring data archival, curation, and tracking capabilities. Different science objectives require configurations of different types, including multiple ensembles of configurations. In this endeavor, we anticipate close collaboration with the HEP community.

Numerically calculating quark propagators requires a mix of capability and capacity computing resources. Nuclear physicists continue to expend significant effort in the efficient generation of propagators; this work, conducted in close collaboration with applied mathematicians and computer scientists, will continue into the exascale era. Calculations in nuclei require further algorithm developments to facilitate increasingly large numbers of quark contractions, including in the area of multi-linear algebra and automated generation. With the growing number of correlation functions that will be produced in the exascale era, significant improvements in data management, including database tools, are essential. Developing these algorithm and codes for exascale architectures will require close collaboration with computer scientists and applied mathematicians. The SciDAC software project, DOE Lattice QCD – Extension-II Hardware Project, and ECP play a critical role in addressing the present-day and future requirements. Early access to exascale architectures through collaboration with vendors is crucial for the nuclear physics community to make optimal use of new computing hardware as it is deployed. With the multistage production process that is inherent to lattice QCD calculations, work flow has become a critical issue.

The structure of hadrons, described by quantities including their mass, charge radius, and parton distributions, plays a crucial role in nuclear physics. Recent precision measurements of the

proton charge radius indicate a significant discrepancy with previous measurements, leading to renewed efforts to measure this basic property. The axial charge of the nucleon is a key element in predictions of the lifetime of nuclei that can undergo $\beta\beta$ -decay, and dictates, in part, the lifetime of a class of radioactive nuclei. The mass of the nucleon is generated by the Higgs mechanism and the quantum fluctuations of the quark and gluons fields, but the precise mechanism of mass generation is presently unknown. Significant effort has gone into exploring the spectra of exotic hadrons, and a major experimental program at JLab — GlueX — is underway to characterize states with important gluonic constituents and exotic multi-quark states. The nuclear physics community has made enormous progress in developing methods to precisely resolve excited hadronic states and to accurately extract scattering information from the relevant lattice QCD calculations. In the longer term, the community's efforts to map out the quantum structures of the proton and nuclei and define their gluonic structure with unprecedented precision would be enhanced by an EIC.

Precise lattice QCD calculations of the properties and reactions of light nuclei, including isospin breaking and electromagnetic effects, will become available with an exascale ecosystem. The results of such calculations will be used, through close collaboration with nuclear many-body theorists, to refine the nuclear force. Present lattice QCD calculations, still at unphysical values of the masses of the quarks, have led to explorations of the chart of nuclides by many-body theorists. Further, theorists are beginning to make both postdictions and predictions of inelastic nuclear reactions using lattice QCD.

The nuclear physics program is currently expanding in the area of fundamental symmetries. To support this experimental program, precision calculations of nuclear systems that have the potential to reveal violations of basic symmetries of nature are needed. With the discovery of neutrino mass, and the possibility of neutrinos violating lepton number, the community is moving toward deploying a ton-scale, $\beta\beta$ -decay experiment(s). Identifying the origins of lepton-number violation, and their importance in nuclei, requires lattice QCD and nuclear many-body calculations. The violation of the combined operation of charge conjugation and parity (CP) symmetry has been known for decades, and general arguments imply that time-reversal (T)-violation should also occur in nature. T-violation has not yet been observed, but searches are underway for electric dipole moments (EDMs) of the neutron (nEDM) and other systems, which would provide direct evidence for T-violation. Significant lattice QCD calculations are required to relate T-violation found in neutrons and nuclei to the structure of fundamental theories.

2.3.5 Hot Quantum Chromodynamics

The study of matter under extreme conditions is required to understand the earliest, hottest moments of the universe and the explosive astrophysical events responsible for producing the heavy elements. To address these conditions, the nuclear physics community requires numerical calculations that employ a number of specialized techniques, including finite-temperature and finite-density lattice QCD calculations and real-time Hamiltonian evolution.

Establishing the phase diagram of QCD is a primary objective of the hot QCD program. Phenomenological modeling indicates that there is a line of first-order phase transitions terminating at a critical point. Closely tied to the beam-energy scan experiments at RHIC, locating the QCD critical point on the phase diagram is the focus of a coordinated experimental and theoretical effort. The theoretical component is coordinated by a topical collaboration (BEST collaboration) that DOE has established. Computationally, the current numerical algorithms used for finite-density lattice QCD calculations require increasingly large resources at each order in an expansion beyond zero density.

The discovery of a near-perfect liquid at RHIC revolutionized our understanding of matter under extreme conditions. Since then, a major experimental thrust has been to characterize this new

form of matter. One probe of the matter is how particles move through it, including correlations between excitations. Transport coefficients and spectral functions provide a direct measure of these features, and the lattice QCD calculations that are essential in characterizing these quantities require extensive capability and capacity HPC resources.

To extract the properties of the QGP and the RHIC liquid by comparing experimental data sets with the predictions of models, scientists must conduct event-by-event processing of the ensembles of QGP initial conditions through relativistic viscous hydrodynamics evolution. These initial conditions are generated by sampling energy depositions by nucleons and their fluctuations through initial state models, such as the Impact Parameter Glasma model. The energy-momentum tensor resulting from this model is extracted from real-time Hamiltonian evolution of the classical Yang-Mills field equations. Including quarks in this evolution is essential for accurate predictions of the early time dynamics of heavy-ion collisions, but this is presently beyond available computing capabilities. Refining these methods to extract the properties of hot QCD matter, by high-dimensional fitting techniques combined with both Euclidean and Minkowski space quantum field theory calculations, will continue into the exascale era.

The need for highly optimized algorithms and codes to perform calculations of QCD at non-zero baryon number density in the exascale era requires closer collaboration between nuclear physicists and computational scientists. In addition to the needs outlined for cold QCD calculations, there are significant additional needs to optimize real-time Hamiltonian evolution that include quantum effects. Porting hydrodynamics codes to heterogeneous architectures is a critical task going forward.

There are significant challenges in the work flow related to sequencing together the distinct types of calculations and simulations that describe heavy-ion collisions in their totality. Combined with the extensive ensemble simulation and high-dimensional fitting requirements, a tightly coupled capability/capacity ecosystem is required. The workforce dedicated to these efforts is currently a limiting factor. Measures need to be taken to avoid lost opportunities resulting from the lack of optimized code for heterogeneous architectures.

2.4 ASCR, NSF, and Nuclear Physics

Collaborations between ASCR, NSF, and nuclear physicists have been key to the successes that the nuclear physics community is currently enjoying. Such collaborations will be essential to accomplishing the scientific objectives in nuclear physics research in the exascale era.

2.4.1 Present Day

Resources provided by ASCR and NSF in the form of compute cycles and workforce, and the ensuing collaborations with applied mathematicians and computer scientists, have been absolutely essential to progress in nuclear physics research and will remain so in the exascale era. An illustration taken from the nuclear physics SciDAC projects, Figure 2-5, shows the interplay between nuclear physics researchers at universities and national laboratories, ASCR researchers, vendors, the SciDAC institutes, and the advanced computing facilities. The left panel is associated with the U.S. Lattice Quantum Chromodynamics Consortium (USQCD) Nuclear Physics SciDAC project (USQCD undated), and the right panel is associated with the Nuclear Computations Low-Energy Initiative (NUCLEI) SciDAC project. The support and organization associated with the SciDAC projects have been critical to forming a coherent effort within the community to develop the algorithms and software elements used in present-day research. Each of these major SciDAC projects is being conducted by a diverse collaboration or team including approximately 120 domain scientists ranging from senior scientists, through postdoctoral fellows, to graduate and undergraduate students; the projects act as vibrant and efficient training grounds for our future workforce.





Figure 2-5. The organization of SciDAC collaborations between NP and ASCR scientists are shown for the USQCD NP project (left panel) and for the NUCLEI project (right panel). (Images reproduced with permission from Robert Edwards and Witek Nazarewicz.)

Capability and capacity computational resources, along with storage and communications, are provided by ASCR, DOE NP, and the NSF. Figure 2-6 shows the computational resources that have been available to two main nuclear physics SciDAC projects. These come from a number of sources. The capability resources are provided by ASCR at its leadership-class facilities and by NSF through Blue Waters computing resources, while capacity computing resources are provided by DOE lattice QCD_extension-II Hardware Project capacity computing hardware at JLab, BNL, and Fermi National Accelerator Laboratory (Fermilab), through allocations at NERSC, university clusters supported by NSF, and through modest Extreme Science and Engineering Discovery Environment (XSEDE) allocations. NERSC also provides significant resources for the experimental program, particularly for data analysis and storage. The proposal- and peer-review-based INCITE (Innovative and Novel Computational Impact on Theory and Experiment) and ALCC programs provide most of the capability computing resources to nuclear physics, including the two major nuclear physics SciDAC projects, and they continue to be critical resources to the field; without them, such research in nuclear physics would stall. In addition, Lawrence Livermore National Laboratory (LLNL) continues to provide substantial computing resources to a third nuclear physics SciDAC project, which is focused on the interface between lattice QCD and nuclear physics.

Cold QCD: Millions of J/ψ core-hrs



INCITE allocations to UNEDF/NUCLEI




Figure 2-6. The capability and capacity computing resources available to the USQCD NP SciDAC project (left panel) and the INCITE awards associated with the NUCLEI SciDAC project (right panel). These do not include resources from ALCC awards or from the NSF. (Image on right reproduced with permission from Joe Carlson; image on left reproduced with permission from Robert Edwards.)

The Major Research Instrumentation program at NSF has provided significant funding for small- to mid-scale capacity computing hardware at universities. These resources have proven to be an important component of the nuclear physics computing ecosystem, because they have permitted the exploration of ideas, algorithms, and codes on short time scales, without the delays encountered in generating proposals that must then be peer reviewed.

The DOE lattice QCD_extension-II Hardware Project continues to be essential for the major advances we have witnessed in lattice QCD during the last several years. Capacity computing hardware and storage are located at Fermilab, JLab, and BNL and leverage the infrastructure of these laboratories. The USQCD Scientific Program Committee annually allocates resources on the basis of submitted proposals, evaluated using multiple criteria, including alignment with the goals of the field and innovative algorithm exploration. In addition, the first large-scale heterogeneous GPU cluster was established by JLab in 2010 and is now part of the DOE lattice QCD_extension-II Hardware Project.

2.4.2 Going Forward

Section 3 presents a significantly more detailed exploration of the status of these major research directions within nuclear physics and specifies the requirements of an exascale ecosystem to optimize the scientific output in each direction. The section also provides estimates of the number of computing cycles, both capability and capacity, and describes how they should be distributed among accelerated heterogeneous architectures and conventional central processing units (CPUs). Storage requirements, from “hot data” (data available for essentially immediate access for a short to an extended period of time) to “cold data” (data that can be retrieved from storage in a reasonable time frame to data archived on tape with the expectation of being touched at intervals of a year or possibly more) are also estimated. The nuclear physics community identified the memory per computing core or accelerator unit for the various directions and found that it spans a wide range. High-speed I/O and data management were found to be critical, and work flow was identified as an emerging issue in large-scale computations. Although it has not been a bottleneck in calculations until recently, work flow is expected to become a focus of development for the exascale era. Similar assessments were made regarding data management, algorithms and codes, curation, and tracking.

Short List of Critical NP Needs

- Enhanced support for a highly skilled HPC workforce within the NP community (domain scientists)
- Enhanced collaborations with computer scientists, applied mathematicians, and statisticians
- Increased floating-point capabilities on conventional CPU and multi-core and GPU architectures to accommodate a mix of projected needs from capability through mid-scale capacity computing resources
- Increased support for capacity computing resources located at universities and nuclear physics laboratories
- Increased memory per computing node with high-speed I/O capabilities
- Large-scale “hot data” and “cold data” storage capabilities
- Rapid-access databases
- Data curation and tracking capabilities, from creation to deletion
- Permanent positions for scientists at the intersection of nuclear physics and applied math/computer science

A skilled workforce is a major component of an exascale ecosystem for nuclear physics. The community identified a clear need for scientists trained in cutting-edge HPC skill sets who — in collaboration with nuclear physicists, computer scientists, applied mathematicians, and statisticians — can jointly tackle the Grand Challenges in nuclear physics using exascale-era infrastructure. An equally important component of an exascale ecosystem, highlighted in *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) and in this exascale review meeting, is the need for local mid-scale-capacity computing resources at universities and laboratories with architectures similar to those of exascale machines. These resources would provide scientists with the following:

- Access to modest computing resources upon demand (without the significant delays imposed by the peer-review process) to test and develop new science ideas, algorithms, work flows, and other features of future large-scale productions.
- A well-established environment and infrastructure in which to educate junior scientists in HPC, moving them from novices with skills related to the use of iPads and similar information appliances; through sequential programming skills, such as are appropriate for cloud computing-scale science problems; to the massive numbers of parallel threads that are required to solve Grand Challenge-scale science problems.
- Platforms with which to explore new configurations of commodity hardware and software.
- Software with enhanced portability across GPU and multi-core architectures.

3 RESEARCH DIRECTIONS AND COMPUTING NEEDS/REQUIREMENTS

3.1 Nuclear Astrophysics

3.1.1 Scientific Challenges and Opportunities

This is a golden era for nuclear astrophysics, a subject that offers tantalizing new opportunities for both nuclear physics and HPC, with the promise of synergistic development of both disciplines. Nuclear astrophysics stands at the intersection of major experimental and theoretical efforts in the areas of nuclear structure and reactions (e.g., FRIB [FRIB undated; FRIB 2012]) and fundamental symmetries (e.g., $0\nu\beta\beta$ decay), and rapidly developing capabilities in astronomical observation — from new ground- and space-based observatories covering the entire electromagnetic band to the advent of gravitational wave astronomy (e.g., LIGO). It is clear then that nuclear astrophysics theory efforts are fundamentally multi-physics in nature, and that large-scale numerical simulations are at the heart of this science. In fact, HPC capabilities are opening up new ways to explore outstanding issues in nuclear physics.

Key objectives in the recent *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) include exploring the origin of the elements (e.g., ascertaining the site[s] where the r-process elements such as uranium, gold, and iodine are synthesized); the physics of ultra-dense neutron-rich nuclear matter; and the nature of neutrinos and their interactions in nuclei and dense matter. Understanding nucleosynthesis means understanding how compact objects work — from core-collapse supernovae (CCSNe), to Type Ia thermonuclear supernovae, to binary neutron star (NSNS) mergers, to explosive phenomena such as x-ray bursts (XRBs). One immediate objective is to leverage the considerable success we have had in modeling neutrino transport and high-temperature (up to tens of million electron volts [MeV]) dense-matter physics in collapsing and exploding compact objects to explore, at a comparable level of fidelity, new venues (e.g., NSNS mergers). Simulations of these mergers could allow LIGO data to provide insight into ultra-dense nuclear EoS physics and determine whether these events are significant contributors to the inventory of r-process heavy nuclei in the Galaxy.

Many simulations in these areas will share some common characteristics: they will be multi-physics and likely multi-dimensional, and they will share a common strategy in which extracting science from the computations will require extensive parameter sensitivity studies. In some cases, unique and challenging new kinds of physics will be modeled — for example, neutrino flavor evolution with neutrino scattering (neutrino flavor/spin quantum kinetics). Progress has already been made in this area, and future efforts will involve a “layered” approach, exploring the efficacy of different approximations. Another example is modeling the electromagnetic signals that might accompany NSNS mergers or the r-process nucleosynthesis of CCSNe, a key part of “multi-messenger” probes of these events.

The multi-physics nature of the HPC techniques employed in this field nicely aligns with the training of adaptable nuclear and computational physicists. Indeed, it is increasingly common to find universities and national laboratories creating joint physics/computing efforts. These efforts attract outstanding graduate students, postdoctoral researchers, and faculty/staff.

3.1.2 Priority Research Directions

Ambitions for the next decade include determining the sites for heavy-element nucleosynthesis and making quantitative predictions for the yields from these sites, as well as predictions for the many observable signals from these events (e.g., gravitational wave and neutrino signatures). The

underlying scenarios and the requisite tools for simulating these explosive events can be grouped into the following concise set of priority research directions, each of which is discussed in the sections that follow:

- CCSNe and NS mergers;
- Neutrino propagation and signals, and dense matter; and
- Type Ia supernovae and XRBs.

3.1.2.1 Core-Collapse Supernovae and Neutron Star Mergers

A major goal on the path to exascale is to understand the mechanism and full phenomenology of CCSNe and the processes that govern the merger of NSNS binaries, including nucleosynthesis and other observables.

Work in recent years has yielded considerable progress toward establishing the explosion mechanism of CCSNe, and much of this progress has been directly attributable to the application of petascale computing resources to simulations of multidimensional radiation hydrodynamics (Lentz et al. 2015; Melson et al. 2015; Bruenn et al. 2016; Takiwaki et al. 2012). Figure 3-1 provides an example of the fidelity of present-day simulations. Although the general picture of neutrino reheating in the context of fluid instabilities has been borne out by these simulations, there is more work to do. Differences in, for example, progenitor structure, incorporated neutrino and nuclear physics, neutrino radiation transport, and the effects of general-relativistic (GR) gravity have led to quantitative and qualitative differences between simulation outcomes. The only way to eliminate the resultant uncertainties is through a series of enhancements to the physical fidelity of all these ingredients. Only exascale resources will allow calculations to be carried out using spectral, multi-angle neutrino transport at high spatial resolutions with state-of-the-art microphysics on evolving GR space times. Furthermore, because CCSNe are believed to occur in stars over a wide range of possible progenitor masses (10 solar masses and above), many simulations will be required to disentangle the effects of massive star evolution from the physics included in the explosion simulations themselves.




Figure 3-1. The specific entropy of the stellar material from a 3-D CCSNe simulation by Lentz et al. (2015).

A similar situation exists in the case of NSNS. Although the last decade has seen a remarkable increase in the sophistication of the nuclear and neutrino physics modeled in NSNS simulations (Rosswog 2013; Roberts et al. 2011; Wanajo et al. 2014), the spatial and temporal resolution requirements for these calculations are considerable. Typically, only a few percent of a solar mass of material is actually ejected in the tidal tails of material produced in the merger, and the neutrino emission is typically confined to less than 10 milliseconds (ms). The sharp differences in density and temperature that evolve during the event also require sophisticated and numerically stable EoSs. All of these effects make the inclusion of multi-physics into these simulations an exascale (and beyond) challenge. Figure 3-2 shows recent results from simulations of NSNS mergers.

CCSNe and NSNS mergers are the primary candidate sites for r-process nucleosynthesis. All of the multi-physics ingredients discussed here impact the viability of the different types of events for the r-process. In particular, the distribution of seed nuclei and the neutron excess necessary to drive r-process nucleosynthesis are sensitive functions of the entropy of the material in the neutron-driven winds that result from CCSNe and NSNS mergers. In addition, the multi-flavor neutrino flux experienced by that material will directly affect the electron fraction (equivalently the neutron-to-proton ratio) of the ejected material. High-fidelity neutrino transport and large nuclear networks are required to calculate these conditions accurately. Furthermore, these components are required at relatively late times (i.e., after the explosion itself has been established). The result is a need for compute-intensive calculations over long, and explicitly evolved, physical times.




Figure 3-2. The collision of two NSs (1.3 and 1.4 solar masses). Shown are volume renderings of the temperature at $t = 1.49, 3.83, 6.27,$ and 8.32 ms after the start of the simulation. (Image available in Rosswog 2013.)

Path to Exascale in Core-Collapse Supernovae and Neutron Star Mergers

- **2016:** Use three-dimensional (3-D) simulations with sophisticated neutrino transport and Lagrangian tracers to calculate r-process nucleosynthesis.
- **2020:** Develop 3-D codes for CCSNe and NSNS mergers with multi-dimensional neutrino transport and large nuclear networks.
- **2025:** Evolve four-dimensional (4-D) space times for CCSNe and NSNS mergers with multi-dimensional transport and r-process networks over a large range of progenitors.

3.1.2.2 Neutrino Propagation and Signals and Dense Matter

Determining how neutrino flavor and spin transport affect both heating and energy/entropy deposition, nucleosynthesis, and potential future Galactic CCSNe neutrino signals will require solving neutrino flavor/spin transport in hot and dense astrophysical environments.

Exploring the frontier of experimental neutrino and fundamental symmetry physics is, in fact, a key goal articulated in the *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015). The following list provides a quick synopsis of where we stand and the outstanding science issues in this field: (1) we know that the three active neutrinos are associated with small masses and, although we do not know what the vacuum mass eigenvalues are, (2) the neutrino mass-squared differences are *measured*; (3) we do not know the neutrino mass hierarchy, whether it is normal or inverted; (4) we do not know the nature of neutrinos (i.e., whether they are Majorana or Dirac in character); (5) we do not know whether there are other kinds of neutrinos with interactions smaller than the weak scale (i.e., sterile states).


We may get a handle on item (1) from cosmological considerations and laboratory experiments (e.g., tritium endpoint, Karlsruhe Tritium Neutrino Experiment [KATRIN], and $0\nu\beta\beta$ decay). The neutrino mass hierarchy (3) is likely to be measured within 5–10 years, making our compact object flavor transformation calculations more predictive for nucleosynthesis and other astrophysical effects. The Majorana or Dirac character of neutrinos (4) will be probed in ton-scale $0\nu\beta\beta$ -decay experiments. Finally, astrophysical considerations, including compact objects, and especially involving active-sterile neutrino flavor transformation and nucleosynthesis, are the only probes of much of sterile neutrino mass-mixing parameter space, especially for larger masses (kilo electron volt [keV] range) and tiny mixing angles.

Much of the key input physics needed for modeling neutrino flavor transformation in compact objects is measured (mass-squared differences and vacuum mixing angles); such measurements set up an exciting and rapidly developing field of computational science: quantum transport.

This new physics may play an important role in the evolution of compact object systems (e.g., in both CCSNe and NSNS mergers). Most of the energy released in these events (as much as ~10% of the *rest mass* of the proto-NS in CCSNe) is radiated as neutrinos on timescales of milliseconds to tens of seconds. The means by which energy and entropy are transported and deposited in these systems is neutrino-flavor-dependent. Boltzmann neutrino transport studies (i.e., without neutrino flavor mixing) show that, at various epochs, the fluxes and/or energy spectra of the different neutrino flavors may differ, so that if one flavor of neutrino transforms into another flavor in the supernova environment, the neutrino heating in the medium can be changed. This process can also affect the neutron-to-proton ratio, a key quantity for nucleosynthesis, especially in the case of the r-process. Key numerical issues in following neutrino oscillations and flavor transformation are

(1) we must follow high-frequency quantum flavor phases, and (2) neutrino flavor transformation in these environments is essentially nonlinear. The nonlinearity stems from the fact that the flavor states of the neutrinos determine how the neutrinos evolve through neutrino-neutrino forward scattering. This forward scattering process and neutrino-forward scattering on charged leptons (principally electrons) together produce flavor-dependent potentials that drive neutrino flavor transformation.

Large-scale simulations of coherent (meaning forward-scattering only) neutrino flavor transformation have revealed that *collective neutrino flavor transformation* can take place (e.g., in a neutrino-heated wind epoch after a CCSNe explosion). An example of such a collective mode of oscillation is shown in Figure 3-3, which illustrates the results of a 3×3 (all three active neutrinos) calculation, providing the survival probability (dark = transformed, light = not transformed) from an initial neutrino flavor, and showing the characteristic “swap” in the normal mass hierarchy (Duan et al. 2010).



The results are dramatically dependent on which of the neutrino mass hierarchies, normal or inverted, nature has chosen. Simulations have revealed that we must compute flavor evolution with a full “multi-angle” geometry if we are to gauge the effects of transformation on r-process nucleosynthesis from neutrino-heated ejecta. This concept is illustrated in Figure 3-4, from the calculations of Duan et al. (2011). The figure shows abundance Y versus nuclear mass number A , along with data (crosses) as labeled.




Figure 3-4. R-process nucleosynthesis including simplified (red line) single-angle calculations of coherent neutrino oscillations, and more complete multi-angle calculations (blue line). (Figure reproduced with permission from Gail McLaughlin; available in Duan et al 2011.)

Coherent flavor evolution simulations along these lines involve solving $\sim 10^7$ to 10^8 nonlinearly coupled Schrodinger-like equations for neutrino flavor. This is a mature subject at this point (Duan et al. 2008; Duan et al. 2010). Although issues remain unresolved (e.g., spatial and temporal instability in coherent flavor transformation), the frontier in this subject now rests on incorporating these calculations into more realistic supernova and NSNS environments. Small density fluctuations associated with turbulence, for example, can also influence neutrino flavor transformation and, in some cases, this effect can be important (Kneller and Volpe 2010).

Another level of complexity in these calculations involves relaxing the coherent, forward scattering-only approximation. Neutrino energy- or direction-changing scattering-induced decoherence can also result in neutrino flavor conversion. The full quantum kinetic equations (QKEs) governing the evolution of neutrino flavor and spin — left- to right-handed (meaning neutrino-to-antineutrino conversion for Majorana type neutrinos) in dense media — have been derived recently (Vlasenko et al. 2014). These QKEs are effectively generalizations of the Boltzmann neutrino transport equations described above, but the QKEs encode flavor mixing and neutrino spin. Solution of these QKEs in a general and complex anisotropic medium (such as that in compact object environments) is daunting; even ordinary Boltzmann neutrino transport that ignores flavor physics is challenging.

However, simple limits of the QKEs have already been treated numerically, with fascinating results. Even a very modest fraction (e.g., 1 in 1,000) of neutrinos suffering a direction-changing scattering can modify flavor transformation, sending quantum mechanical flavor information inward, toward the NS(s) source. Such scattering changes the otherwise coherent neutrino flavor transformation problem from an initial value problem, with neutrino flavors specified on the neutrino sphere, into something more akin to a boundary value problem. This phenomenon is called the neutrino “halo” effect (Cherry et al. 2012). The halo has been successfully simulated for oxygen-neon-magnesium

(O-Ne-Mg) CCSNe. A goal of current simulations is to extend such a treatment of the halo effect to iron (Fe)-CCSNe. Figure 3-5 indicates the surprising potential leverage of the neutrino halo: the left panel provides a color-coded density plot of a two-dimensional (2-D) model, and the right panel provides an indication of the neutrino direction-changing alteration in the potentials that govern neutrino flavor conversion (Cherry et al. 2012).




Figure 3-5. Potential impact of neutrino-matter interactions in CCSNe — the so-called ‘halo’ effect. (Image reproduced with permission from George Fuller; available in Cherry et al. 2012.)

QKE calculations are a frontier issue in modeling early-universe physics at the weak decoupling epoch. The high degree of symmetry in this environment reduces the complexity of QKE calculations, allowing simulations to be performed using existing supercomputing capabilities. Progress on these spatially homogeneous and isotropic conditions may provide insights into how to take better advantage of the increased memory and speed of exascale supercomputers to explore QKE physics in compact objects.


Finally, QKE studies have revealed that neutrinos propagating coherently in an anisotropic medium of matter and neutrinos are in coherent superpositions of flavor states and helicity (spin) (i.e., left- versus right-handed) states (Cirigliano et al. 2015). Just as flavor can be transformed, so spin can be transformed coherently, albeit with daunting adiabaticity criteria. Of course, in the case of neutrinos of Majorana character, this process is tantamount to transforming neutrinos into antineutrinos and vice versa, a process to which the local neutron-to-proton ratio can be highly sensitive. Interestingly, the spin conversion is sensitive to absolute neutrino masses and Majorana phases; the conversion accesses physics that is complementary to that probed in proposed $0\nu\beta\beta$ -decay experiments.

Path to Exascale in Neutrino Propagation and Signals and Dense Matter

- 2016: Simulate *coherent* neutrino flavor transformation using realistic supernovae models.
- 2020: Obtain multi-dimensional coherent neutrino flavor solutions; treatment of the “halo” for Fe-core collapse models. Obtain full neutrino quantum kinetics for homogeneous and isotropic (early universe) environments with full coupling to nuclear reactions.
- 2025: Implement approximate quantum kinetics for flavor/spin in CCSNe and binary NSNS merger environments.

3.1.2.3 Type Ia Supernovae and X-ray Bursts: Progenitors, Explosion Mechanisms, and Nucleosynthetic Yields

Type Ia supernovae (SNe Ia) are thermonuclear explosions of carbon/oxygen (C/O) white dwarfs (WDs) in binary systems (Figure 3-6). They are bright enough to be visible over a considerable fraction of the observable universe. This fact led directly to their integral role as distance indicators in the discovery of the accelerating expansion of the Universe (Perlmutter 2012). Although observations have helped to constrain many of the properties of SNe Ia, a persistent question remains: what is the progenitor? Possibilities include a single WD accreting from a companion star (single degenerate scenario) or two WDs that inspiral and merge or collide (the double degenerate scenario) (see, for example, Calder et al. 2013). No conclusive observation of the progenitor system exists, but in all cases, the majority of the C and O in the WD(s) is converted into iron/nickel (Fe/Ni) and intermediate-mass elements such as silicon (Si), and the liberated nuclear energy release unbinds the star. For many years, the Chandrasekhar-mass single degenerate model (henceforth, the Chandra model) received the most attention. The Chandrasekhar mass is the maximum possible mass of a WD, so explosions at this mass imply that SNe Ia would be alike. But contemporary observations suggest a measure of diversity. In addition, massive C/O WDs in binary systems are rare, leading to questions of the rate of events from this scenario. Accordingly, research in recent years has diversified as well, now encompassing virtually all explosion models. Topics of present interest include merging WDs, sub-Chandra explosion models, and the convective Urca process that is believed to occur during the late stage of accretion in the Chandra model.



The standard picture for mergers has the WDs inspiraling as gravitational radiation removes orbital energy. The less massive star becomes tidally disrupted and the more massive star accretes this material. A longstanding concern is that when this mass transfer begins, thermonuclear burning can ignite at the edge of the star, converting it to O/Ne/Mg and leading to the collapse of the WD into an NS, instead of giving rise to a SNe Ia. This system is inherently 3-D, and only through detailed simulation can we understand the dynamics of the mass transfer, and thereby assess the feasibility of this model given these sensitivities.

Sub-Chandra models have the advantage that systems that contain a moderate-mass WD ($0.8\text{--}1.0 M_{\odot}$ [solar masses]) are abundant, and these models can also reproduce the known delay-time distributions and SNe Ia rates — something the Chandra model struggles with (Ruiter et al. 2011). In these models, the WD accretes a layer of helium (He) on its surface that detonates, sending a compression wave into the underlying C/O WD and igniting a second detonation that unbinds the star. Understanding and producing quantitative predictions will require a realistic 3-D convective field in simulations, as well as accurate modeling of the subsequent evolution of the detonation. Ultimately, large ensembles of all possible explosion scenarios will be required to establish the progenitor site(s); each of these ensembles will require high spatial resolution with large nuclear networks — both to achieve adequate fidelity in energy generation and to establish nucleosynthetic yields that can be used to confront observations.

XRBs are the thermonuclear runaway of a thin hydrogen/helium (H/He) layer on the surface of a NS. This fuel layer, accreted from a companion star, is compressed to the point of ignition by the immense gravity of the NS. One-dimensional (1-D) hydrodynamic studies can reproduce many of the observable features of XRBs: burst energies ($\sim 10^{39}$ erg), rise times (seconds), durations (tens to hundreds of seconds) and recurrence times (hours to days). However, 1-D models assume that the fuel is burned uniformly over the surface of the star, which disagrees with observations of oscillations in the rise of the light curves of XRBs associated with the rotation-modulated spreading of a burning front across the NS. Before the actual outburst, the burning at the base of the ignition column drives convection throughout the overlying layers and sets the state of the material in which the burning front propagates. A proper treatment of the convection in these extreme conditions, free from parameterizations, requires 3-D simulations with large nuclear networks.

Importantly, XRBs are sensitive probes of the nuclear EoS and rp-process nucleosynthesis. Observations using the National Aeronautics and Space Administration's (NASA's) Rossi X-ray Timing Explorer (RXTE) have resulted in the discovery of fast (200–600 Hz), coherent X-ray intensity oscillations in XRBs (Strohmayer 2004) (Figure 3-7). These oscillations result from the aforementioned spin modulation of the thermonuclear burst flux from the NS surface, which is determined by the radius and spin rate of the underlying NS. Because of this relation, detailed modeling of burst oscillations can be an extremely powerful probe of NS structure, and thus the EoS of supra-nuclear density matter. Also related to structures in the burst profile, the effects of hydrodynamic instabilities during the thermonuclear explosion must be calculated accurately to differentiate those effects from other effects (e.g., uncertainties in the masses of neutron-deficient isotopes and proton capture rates). Finally, establishing the total time of the outburst is essential to determining the endpoint of rp-process nucleosynthesis. All of these goals require reactive hydrodynamics calculations be carried out in the strong gravity of the NS surface.




Figure 3-7. Vertical velocity showing the convective structure in a low-mach XRB simulation. (Image available in Zingale et al. 2015.)

Path to Exascale in Type 1a Supernovae and X-Ray Bursts

- 2016: Use 2- and 3-D calculations of SNe Ia and XRB to determine features of reactive flow in the events and the potential impact on nucleosynthesis.
- 2020: Identify the progenitor(s) of SNe Ia and make robust predictions of light curves and nucleosynthesis.
- 2025: Calculate flame spreading across the NS surface in XRBs, including robust subgrid flame model and coupling to photon transport.

3.1.3 Cross-Cutting Research Directions

Development of hardware such as GPUs requires that issues of data locality and movement be addressed, for both new developments — such as hierarchical memory — and more established approaches such as threading. For complex multi-physics codes, particularly codes that rely on other software packages (such as adaptive mesh refinement [AMR] libraries), the development task may be daunting and involve many person-years of effort. Meaningfully utilizing threading, for example, may require a data layout that is at odds with both extant data storage in a mesh package and the stencil of a given operator. The optimal data layout may well be hardware-dependent. Finding the right compromise between performance and portability is itself a research topic, and implementation may require platform-dependent development via approaches such as frameworks in which data storage and locality are abstracted from the application developer.

Traditionally, collaborations to develop large simulation codes have comprised teams of computer scientists and application scientists working together. Fields evolve, however, and this model is now less common. Increasingly, expertise in both HPC and algorithm development resides not in Computer Science departments but in Applied Mathematics and Computational Science departments and among the application scientists themselves. The issue is particularly acute in the education and training of students. NP and ASCR education and development efforts should accordingly provide support for a wider spectrum of investigators.

Nuclear astrophysics graduate students are trained in computational techniques and can make valuable contributions outside of the field. This is one important reason why the field is successful in attracting excellent students. Indeed, demand exceeds the number of students currently available. Nuclear astrophysics groups could easily handle more students if more support were available. There are strong demands for graduates trained as nuclear astrophysicists at the national laboratories (including NNSA laboratories) and DOE computing centers. Often, even if these physicists switch to computing support roles or take on programmatic tasks at the laboratories, they keep their ties to nuclear astrophysics and continue to make contributions to the field. The graduates are also desired by industry because HPC and data science are in-demand skills at tech companies.

3.1.4 Computing, Data, and Networking Needs and Requirements

Effective use of exascale systems for nuclear astrophysics simulation will involve requirements spanning the ecosystem — from the wide-area network connections provided by ESNet to the individual processing cores on each node. Many of these requirements are driven by the need to perform ensembles of simulations to achieve important scientific results. Dozens to hundreds of individual simulations of the stellar explosive events described in Section 3.1.2 (Priority Research Directions) will be required. The intent is to fully explore the space of possible progenitors, initial system conditions, and necessary parameterizations in each scenario. In addition, the impact of new or newly implemented (and more complete) physics at each epoch along the road to exascale will require a re-evaluation of earlier work.

Because much of nuclear astrophysical simulation is multiphysics in nature, possibilities for collaboration with ASCR-supported researchers in algorithm development for multiphysics abound. Applied mathematicians and computer scientists with interests in, for example, methods to move away from operator splitting for coupling physical models, would be excellent candidates for collaboration. Work in other areas — ranging from dense and sparse linear solvers to high-order techniques for transport to advanced AMR techniques — is also needed.

In this context, important features of the exascale ecosystem required for nuclear astrophysics calculations include the following: wide-area networking, disk and archival resources, burst buffers, internode networks, memory hierarchy, and processor and programming models.

The number of reduced-data products that could be useful to a broad audience and could be made publicly available from nuclear astrophysics simulations is relatively small compared with, for instance, the number from large-scale structure simulation or from terrestrial plasma physics simulations. We anticipate that total data volumes will not exceed several terabytes (TBs) by 2020 and perhaps 10–50 TB by 2025. These data products include simulated spectra from reactive and radiation hydrodynamics simulations, as well as multi-messenger signals, such as neutrino signatures and gravitational wave templates. Web-based interfaces to access, visualize, and interact with stored data would be useful for science analyses and for disseminating the spectral output to the community. Given the limited number of full simulations that will be available, it would be helpful to field statistical tools (e.g., emulators based on Gaussian processes) that interpolate over the populated parameter space and find best fits to observational data. However, most of the investigators in the field perform related simulations on a multitude of platforms, both ASCR

facilities and NSF machines (and, in some cases, on NASA resources as well). Because the size of individual checkpoint files can grow to be large (see below), being able to efficiently transfer these files over the wide-area network becomes a challenge.

Data storage capacity is a marked need. Sheer capacity is a necessary, but not sufficient, condition to enable science from nuclear astrophysics simulations. The spatial resolution requirements and multiphysics nature of nuclear astrophysics simulation lead directly to this need for considerable storage. For example, the largest, highest-fidelity CCSNe simulations today produce roughly 500 TB per second of physical evolution. Increasing local spatial resolutions (via AMR) by two to three orders of magnitude, coupled with an increase of 10–100 $\times$ in the number of degrees of freedom at each grid point (e.g., from fully angle-dependent neutrino transport and/or from increasing the number of nuclear species evolved) will lead to data volumes of roughly 100 petabytes (PB) per simulation by 2025. In addition to the immediate requirement for capacity for long-running simulations, these data need to have relatively long persistence on resources that can enable interactive data analysis; for example, sufficient read bandwidth and low latencies for read. The timescale associated with the peer-reviewed journal publication cycle—including addressing the comments and questions of referees—is the minimum persistence time required. Although increased physical fidelity will require significant additional computational load, it should be noted that although data volumes are expected to increase significantly in the exascale environment, the rate at which data will be pushed to disk will not increase as quickly. Therefore, capacity requirements far outpace I/O bandwidth requirements for the multiphysics simulations that characterize nuclear astrophysics.

Nuclear astrophysics simulation codes already universally make extensive use of checkpoint/restart. Indeed, considerable code engineering has been performed in many cases to ensure that maximum bandwidth/minimum write times are achieved. Ideally, the addition of burst buffer technology in future platforms would present use methods that would allow this significant investment in code to persist. Although small additions or changes to application programming interfaces (APIs) would present few problems, these changes would hopefully be portable and are necessary to achieve good performance. An alternate burst buffer use case is using non-volatile random-access memory (NVRAM) as out-of-core storage for large, multi-dimensional data sets during runtime. Typical uses of such an out-of-core capability would be storage and read of tabular physics, such as neutrino and photon interaction cross-sections and thermonuclear reaction rates. These data sets would be write-once (at runtime), read-many in nature, and large enough to exceed single-node memories.

More bandwidth is also needed between the new, more powerful nodes that will become available. The combination of powerful nodes and high-fidelity multiphysics will lead to large potential communication payloads, perhaps as much an order of magnitude larger than the nearest-neighbor communications performed today. Latency is also an important consideration, but the requirements for global reductions are less stringent and generally do not grow with increased physical fidelity (e.g., global timestep determination scales with physical resolution, at worst).

Deep memory hierarchies will present significant challenges for multiphysics nuclear astrophysics simulation. Because large, multidimensional data structures are ubiquitous in astrophysics implementations, there are issues with both memory layout—how data structures are arranged in memory to ensure that memory can be accessed efficiently (e.g. for vectorization)—and memory placement. Memory placement issues include placing data structures in particular tiers of the memory hierarchy to allow the best locality (e.g., in GPU memory if the operations are to be performed by the GPU). Most importantly, the nuclear physics community needs a portable way to efficiently handle both memory layout and memory placement. Libraries are likely the best method because they do not require direct compiler support (helping with ubiquity) and can be “carried around” with individual codes, if need be. Completely ceding control of memory placement and layout to the runtime system is not considered viable or desirable because layout and placement needs are physically motivated and can be time dependent.

To most effectively use the various processing elements that will be available at the exascale and in the era leading up to it, a more-or-less portable programming model for node-level parallel processing is desirable. Several groups have already adopted a more descriptive approach to this problem (e.g., through the use of OpenACC), but more prescriptive approaches are also used. Whatever the nature of the exascale programming model, reasonable default assumptions about data layout, data placement, execution, and dependencies will be the distinguishing characteristics of a successful model. Wide availability of compilers and runtime systems that implement this model will be key to its ultimate adoption by the community. Significant amounts of development on smaller resources — from laptops to local clusters — will be required to produce efficient codes for the exascale. Students and postdocs at universities will undertake much of this development, tying the need for wide availability directly to workforce development.

Most current nuclear astrophysics codes are written primarily in Fortran, and the continued support of this language is important for most groups. Ultimately, it is the robust support of arrays in Fortran that make it the preferred language for most of the codes in the field, regardless of the historical reasons for its use. Nevertheless, although C++ is becoming more and more popular as a “framework” language for nuclear astrophysics, there are decades of development encapsulated in hundreds of thousands of lines of Fortran in the field.

Figure 3-8 summarizes the computing resources required to accomplish the planned nuclear astrophysics program. In all computational resource dimensions, nuclear astrophysics has significant exascale computing requirements.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025 NUCLEAR ASTROPHYSICS




Figure 3-8. The capability and capacity computing resource requirements (left panel) and the hot and cold data requirements (right panel) in 2025 to accomplish the science objectives of the Nuclear Astrophysics program.

3.2 Experiment and Data

3.2.1 Scientific Challenges and Opportunities

The scientific challenges and opportunities associated with nuclear physics Experiment and Data span the full science program of DOE's NP program. These challenges and opportunities were well articulated in *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) and include the following:

- Unfolding the quark and gluon structure of hadrons and nuclei;
- Realizing a predictive model of nuclei and their role in the cosmos;
- Testing the particle–antiparticle nature of neutrinos and other fundamental symmetries and neutrino research that open new doors to physics beyond the Standard Model (BSM); and
- Studying the properties and phases of quark and gluon matter in the high temperatures of the early universe, and exploring the spin structure of the proton.

Rather than pushing the energy frontier, nuclear physics experiments are driven by precision to access the multi-dimensional and multi-channel problem space, requiring beam intensity, polarization, and exquisite control of backgrounds and systematics. Over the next decade, this science will be addressed at the current and planned experimental facilities in the nuclear physics portfolio. NP supports three national user facilities:

1. The Argonne Tandem Linac Accelerator System (ATLAS) accelerator at Argonne National Laboratory (Argonne) enables experiments in low-energy nuclear structure and astrophysics.
2. Science at the Continuous Electron Beam Accelerator Facility (CEBAF) at JLab focuses on nucleon structure and QCD.
3. The RHIC at BNL is used to study the properties of the QGP and the spin structure on the nucleon.

The nuclear physics community also uses the Fundamental Neutron Physics Beam Line at Oak Ridge National Laboratory (ORNL) to perform precision studies of the neutron and its interactions. FRIB, under construction at MSU, will produce the short-lived isotopes needed to explore nuclear structure, nuclear astrophysics, and fundamental symmetries. NP is planning to deploy a ton-scale instrument to explore $0\nu\beta\beta$ decay, providing a sensitive search for physics BSM that leverages current efforts such as the Cryogenic Underground Observatory for Rare Events (CUORE) experiment depicted in Figure 3-9. Beyond the next decade, there has been an expressed need from the community for construction of a new powerful EIC to precisely elucidate the role of gluons in nucleons and nuclei.

The successful realization of science outcomes at these facilities critically depends on the optimization of experimental infrastructure. HPC plays an important role in the reliable and efficient execution of the experiment and in the rapid and well-grounded treatment of the resulting data.




Figure 3-9. The 750-kilogram (kg) tellurium dioxide (TeO₂) bolometer array assembled by the CUORE Collaboration to search for the $0\nu\beta\beta$ -decay mode of nuclei. (Image reproduced with permission from Yury Suvorov [UCLA, LNGS] and the CUORE Collaboration.)

The particle accelerators at the heart of the nuclear physics user facilities rely on computationally intensive simulations to optimize accelerator performance. These include simulations of the beam transport through the machine, including the self-field interaction within the beam itself and the beam–beam interaction from the opposite colliding beam. Design of the future EIC will offer significant computational challenges. There is a near-term need to assess competing accelerator designs for technical pitfalls and feasibility, as well as a longer term need for a full-fledged design study to optimize cost/performance. Both these needs can only be met with HPC, because the self-field effects and the electron cooling that need to be considered for the EIC design are based on self-consistent particle simulations that demand exascale computing resources.

Reliable and safe accelerator operations performance has elements that benefit from HPC. The overhead associated with the tuning of a complex particle accelerator can be significantly reduced, and the performance of the accelerator can be significantly improved, by employing machine learning and real-time multi-physics beam transport simulations, both of which are computationally intensive. The high-power particle beams can produce prompt radiation, and verification of the shielding for personnel protection makes use of radiation transport codes that demand significant computational resources to achieve results that have the necessary statistical significance. The damage caused by the interaction of high-power beams on targets and beam dumps can severely limit facility operations. Radiation transport and commercial engineering codes are both used to simulate material properties under such severe conditions, where again significant computational resources are needed for the requisite statistics.

Highly efficient detector systems are necessary to achieve the science goals outlined above, especially when the underlying experiments aim to detect rare events. Detector simulations are performed by using standard codes such as Geant, Fluka, and virtual Monte Carlo (VMC). As detector complexity has increased, the compute needs for the simulations have grown correspondingly. Of course, the detector simulation itself is just one part of the work flow for simulated and observed event reconstruction. Event generators follow different paradigms depending on the physics models involved. Their computational needs vary, but the demand for

statistically significant results typically translates into computational requirements that comprise more than half of the total computational resources needed for experimental data analysis.

The approach to the acquisition of data from the highly efficient detector systems has not changed much in the last 20 years. Detector signals are digitized by electronics in front-end crates, and data readout is initiated by the trigger electronics. Data are transported to an event builder, and the subsequent built events are distributed for filtering, monitoring, display, etc. The event stream is also stored to disk for offline analysis. The trigger defines the data; therefore, experiment optimization depends critically on minimizing the time needed to validate data. Data throughput is determined by the physics event rate, detector response, and overhead associated with the data work-flow steps enumerated above. The latter depends on the performance of both the network and compute environment. A current trend is to move some functionality previously performed in software running on embedded processors into firmware on custom electronics. The ever-increasing expense and decreasing availability of beam time at user facilities drive such efforts. Loosening triggers to store as much data as possible for future data-mining, without putting the primary science goal of the experiment at risk, is becoming more commonplace.

In terms of data analysis and data-mining, there are a plethora of ways that data are treated by the different experiments. NP-supported experimentalists generally use common software building blocks (e.g., Geant, ROOT) for modeling and analysis of data. In general, HPC methods and advanced computing facilities are not being utilized to process the data. The compute infrastructure for data analysis and data storage/archiving usually resides within the host facility. The offline approach to data analysis is usually different from that performed online, which as stated earlier is “limited” by the need for rapid data validation. However, there is already motion in the field to employ higher performance online computing environments to narrow this gap between online and offline analyses. The nuclear physics experiment community should continue to explore other opportunities to leverage HPC resources to optimize data analysis, while still being able to employ established tool sets.

3.2.2 Priority Research Directions

The research directions of the nuclear physics experiment community cover the full breadth of scientific study supported by the NP research program. There are common methods and tools needed by nuclear physics experiments to acquire, process, and analyze future data from DOE NP facilities. These methods and tools are split into three categories, which map to three broad (and sometimes overlapping) areas of potential collaborative research with ASCR scientists:

- Data streaming and near-line validation at NP facilities;
- Detector simulation and data analysis on HPC resources;
- Complex work flows, data management, and network infrastructure; and
- Accelerator simulation.

Each of these is addressed in the sections that follow.

3.2.2.1 Data Streaming and Near-line Validation

In almost all current nuclear physics experiments, data acquisition systems use an event-oriented, triggered readout model. Detector hardware and custom electronics are used to provide a trigger signal to identify when an interaction has occurred in the detector. The trigger electronics define a time window within which digitized data are part of the interaction or event. The trigger makes a (fast) decision on whether to continue digitizing the event or discard it. As data from slower

detectors become available, further trigger decisions are made, and events surviving the trigger filter are read from the digitizing electronics. In some cases, the events are further filtered by software to reduce the data rate to permanent storage.

Although this model has been very successful, it has several inherent limitations. Because the data are partitioned into events by the trigger electronics, such systems are limited to interaction rates that guarantee that the time window defined by the trigger contains data from only one event. Typically, overlapping events are discarded. Although the first stages of the trigger are implemented by electronics (e.g., firmware in application-specific integrated circuits [ASICs]), complex trigger decisions take a finite time. Elaborate trigger schemes have been implemented to reduce the dead time when the trigger cannot process a new event because it is processing the previous one. Triggered schemes always entail the inherent cost of discarding potentially useful data — either through misconfiguration or through design choices issuing from the compromise between trigger accuracy and speed.

An alternative data acquisition model, being pursued for instance, for the GRETA experiment to be deployed at FRIB (see Figures 3-10 and 3-11), is streaming readout. In this model, the digitizing electronics produce a time-stamped stream of data. The combined data streams from all of the detectors are processed by a combination of firmware and software in real time. The limitation of such a model is that, for high-luminosity experiments, the data rates involved are large enough that it is difficult to affordably implement such a system using current technology. This will almost certainly not be the case on the 5- to 10-year time scale. There are several advantages of such a system. Because the data streams are not partitioned into events, algorithms can be implemented that recognize events within the stream and untangle overlapping events, allowing much-higher-luminosity experiments than currently feasible. The data streams could be written to temporary storage. The streamed data would persist long enough for online validation, while a software filter could be applied to reduce the rate to permanent storage. It would be possible to run more than one filter algorithm at the same time to examine different aspects of the same data or to test new algorithms on live data. A streaming readout system requires a combination of HPC and storage coupled to the detector by a low-latency, high-bandwidth network. A software framework to implement a data acquisition system in such an environment is also required and should include data transport, messaging, system integration, and visualization components.




Figure 3-10. GRETA is the predecessor for the GRETA gamma-ray tracking detector array. The construction of GRETA, a full 4π coverage device, with increased rate requirements, coupled with the need for real-time processing, requires significant I/O, CPU, memory, and storage resources. (Image by Kevin Bailey of Argonne National Laboratory; reproduced with permission from Shaofei Zhu.)




Figure 3-11. The High-Resolution Spectrograph (HRS) is one of the planned spectrometers for FRIB. Although the rates of rare isotope experiments may be small, such experiments face data-handling challenges with multiple data stream merging and the treatment of online, near-line, and offline analysis. (Image by Erin O'Donnell, Thomas Baumann, Don Lawton, and Ben Arend and reproduced with permission from Remco Zegers; available in Gade and Zegers 2014.)

Path to Exascale in Data Streaming and Near-line Validation at NP Facilities

- **2016:** Conduct collaborative research into high-precision, real-time analysis of streamed data at DOE experiment facilities.
- **2020:** Develop common tool sets for using HPC-like facilities at experiment sites to perform high-ingest data analysis and validation tasks on triggerless data streams.
- **2025:** Implement end-to-end software frameworks that employ best practices for integrating data transport, messaging, compute services, and visualization components.

3.2.2.2 Detector Simulation and Data Analysis on HPC Resources

In contrast to the processing of experimental data, in which the primary throughput bottleneck is data movement between memory and storage, detector simulation is the CPU-intensive aspect of experimental nuclear physics. Up to 70% of the total CPU time available to an experiment for offline computing could be spent on detector simulations. Nearly all nuclear physics experiments make use of a common toolkit for detector simulations called Geant, developed and primarily supported by the HEP community. The current production version, called Geant4, supports multi-threaded parallelism at the event level, but the package is not suitable in its present form for deployment on petascale-type HPC resources because of the size of the data structures required to hold an entire event in memory. Efforts to develop a vectorized detector simulation package (called Geant V) are ongoing. In the meantime, it may be possible to adapt the Geant4 framework to incorporate coarse-grained (subdetector and track-based) parallelism and, in this way, reduce the memory requirements per thread by perhaps as much as an order of magnitude from the present 1–2 gigabytes (GB). Improvements to Geant4 to use HPC resources more effectively will be of significant benefit to the community in the short term. For the longer term, supporting projects such as Geant V will greatly benefit both the nuclear physics and the HEP communities.

There is a transitional issue associated with most experimental analysis workloads. Those workloads are parallelized by “events.” An event is an intellectually natural unit, corresponding to a discrete time slice intended to capture a physical occurrence or interaction as it unfolds in the detector, such as the complex relativistic heavy-ion collision depicted in Figure 3-12. Using this natural parallelism for processing has been advantageous for many years and is part of the fundamental structure of some of the widely used community software tools such as ROOT, used in production for more than 20 years as analysis frameworks. Benefitting from emerging architectures will require different approaches. At present, individual nuclear physics collaborations develop their own software frameworks for data analysis, typically incorporating components from ROOT but deploying them in different ways, guided by habits and practices that have been built by experience within the nuclear physics community. For the longer term, it will be beneficial to look beyond in-house HEP/NP solutions for analysis tools and frameworks. In addition to introducing parallelization for data mining on relatively large data sets, it could also facilitate the use of standardized statistical and advanced visualization packages




Figure 3-12. An ALICE event display of a real Pb-Pb collision at 5.02-tera-electron-volt (TeV) nucleon center of mass energy. (Image reproduced with permission from European Organization for Nuclear Resources [CERN]; available on the CERN document server.)

Path to Exascale in Detector Simulations and Data Analysis on HPC Resources

- 2016: Conduct collaborative research to adapt Geant4 simulation and ROOT-based data analysis frameworks for coarse-grained parallelism on HPC.
- 2020: Develop vectorized detector simulations (Geant V) on HPC and models or best practices for data analysis on emerging architectures.
- 2025: Develop new software frameworks for fine-grained parallelism for detector simulations and data analysis.

3.2.2.3 Complex Work Flows, Data Management, and Network Infrastructures

Data processing in nuclear physics experiments spans a diverse range of activities, from raw data reduction into physics-usable quantities to the large variety of end-user analyses. Complex work flows are often required to link numerous computing tasks that typically manage one or more tasks — data access, data movement, data storage, or processing provenance — while preserving the integrity of the scientific results. The work flows vary greatly in scope (local, distributed, across heterogeneous resources), level of automation and abstraction, and adaptive

capability (e.g., responding to network clutter or resource availability). Not surprisingly, individual experiments have come up with a plethora of in-house solutions implementing different levels of sophistication. Lack of support for collaborative or integration efforts has resulted in sparse sharing of knowledge across NP, leading to unnecessary duplication of work with varying results. As experiments become more complex and computing landscapes become more diverse, we see an opportunity for collaborative research and development between NP, ASCR, and NSF that supports scientists in building and integrating tool sets that can be reused by many experiments and incorporated into both experiment site operations and ASCR-supported facilities.

An illustrative example is a design trend at JLab, the Large Hadron Collider (LHC), and RHIC in which tasks are broken down and run as detached processes; job steering and data are accessed through fault-tolerant message queues and/or steered by global “reasoners” that optimize every element of the ecosystem as a resource. Those techniques have been shown to be adaptive, scalable, and resilient, but they have not propagated to the wider NP community. The implementations imply a high level of network communication and information flow, and the approach begins to resemble patterns referred to as an “internet of things.” Such a radical paradigm shift would require an architecture dialogue between domain scientists, facilities, and ASCR and NSF computer scientists.

The research potential is broad, encompassing online tasks for detector controls and data collection, processing, and filtering; offline pattern recognition jobs to construct interpretable physical entities; and optimized end-user analysis. To meet this challenge, we suggest the creation of collaborative programs aimed at consolidating knowledge and producing common tool sets that provide automated, distributed, secured, fault-tolerant, and dynamically adaptive work flows to be implemented and deployed at many scales: from facilities at the experiment sites, on distributed institutional clusters, and on leadership-class facilities operated by ASCR and NSF.

Data processing at the petascale level on distributed systems presents unique challenges. Robust wide-area-network connectivity and capacity (tens to hundreds of GB per second) are needed between site borders, while recent evolutions in data access also require significant functionalities such as the distribution of software and environments all the way to node-level processing. Beyond access and transfer problems, work flows spanning security domains are possible only with the existence of and reliance on secured interfaces, federated identities, and common authentication protocols across sites, not only to access computing resources but also to appropriately share data among collaborators and other authorized researchers.

The Open Science Grid (OSG) has addressed many of those challenges, with active network monitoring collaborations with ESNet and Internet2, software distribution with OASIS (CERN VM File System [CVMFS]), and federated identities using X509 certification now under the CILogon CA (<http://ca.cilogon.org>). While they are the cornerstones of distributed computing in NP, the tools, interfaces, and solutions developed by the OSG have not been generally available at the larger ASCR facilities. Rectifying this disparity represents a significant opportunity for collaborative efforts that will enhance NP experiment productivity by making ASCR flagship resources accessible to the experiment community.

Path to Exascale in Complex Work Flows, Data Management, and Network Infrastructure

- **2016:** Conduct collaborative research with NP scientists to develop adaptive, fault-tolerant data-processing work flows that optimize resources.
- **2020:** Develop common-use work-flow tool set that connects experiment site operations, institutional clusters, and ASCR-supported facilities.
- **2025:** Create end-to-end software frameworks that operate at different data-processing stages to efficiently use available computing resources and adapt to real-time environments.

3.2.2.4 Accelerator Simulation

Accelerator simulation for the next-generation NP accelerator, for example, the EIC, is computationally intensive. Beam–beam forces from the electromagnetic fields of the opposite colliding beam during the collision and nonlinear space-charge forces from the Coulomb interaction among charged particles at lower energy can drive beam emittance growth and particle losses and limit the beam bunch intensity and the collider luminosity. In order to self-consistently model these effects, one needs to solve multiple Vlasov-Poisson equations. This is done using a particle-in-cell–based method, in which particles and fields are advanced self-consistently. Meanwhile, electron cooling (coherent and bunched beam) is proposed to counter the growth of beam emittance in the EIC design. These cooling schemes are also modeled using a particle-based method. A number of tools are available to model these effects. The BeamBeam3D is a massive parallel particle-in-cell code to model beam–beam effects on petascale supercomputers. The GHOST is a particle-based code to model beam–beam effects on GPUs. The IMPACT code suite is a massive parallel particle-in-cell code to model space-charge effects on petascale supercomputers. BETACOOOL and VORPAL were used to model electron cooling. All these particle-based codes require significant computational resources during a simulation. Global, across-processor communication is used in the particle-in-cell codes. A high-bandwidth, low-latency network connection in the next-generation supercomputers will enhance the performance of these codes.

3.2.3 Cross-Cutting Directions

3.2.3.1 Nuclear Physics Community

Experiment and Data have crosscuts across all areas of NP. The science opportunities in Experiment and Data, as noted previously, span the full science program — including nuclear structure and reactions, cold and hot QCD, nuclear astrophysics, and fundamental symmetries. There is strong synergy between experiment and theory in all these areas. As an example, sensitivity studies carried out using computationally demanding explosive nucleosynthesis reaction network calculations identify key nuclei whose nuclear properties are then determined at the relevant NP national user facility. Another example is the modeling of the reaction dynamics of relativistic heavy-ion collisions in multiple dimensions. This modeling is compute intensive and is important for exploring aspects of hot QCD and the properties of the QGP, as well as serving as the event generator in the first step of work flow for detector simulations. These two examples give a flavor for the overlap of experimental and theoretical efforts within NP. Theory researchers in NP now rely heavily on HPC resources to realize scientific outcomes. Overlaps between theory and experiment are expected to grow as HPC becomes more commonplace in the experimental realm.

NP Experiment and Data also has HPC crosscuts with activities supported by HEP and BES within DOE SC. Both HEP and BES have accelerator-based national user facilities. The calculations and simulations for accelerator design and performance optimization — which, as discussed previously, are computationally demanding — are pertinent to the full range of energies and ion species across NP, HEP, and BES. Operations at all the accelerator facilities in SC would benefit from automated tuning routines, which require the implementation of machine learning and real-time multi-physics beam transport simulations, both of which are computationally intensive. NP also shares with HEP and BES the need for advanced approaches to data acquisition, robust sharing of large data sets among remote collaborators, and storage and archiving of data to meet data management requirements. NP and HEP experiments operate a number of similar scoped, multi-component detector systems with complex triggers that require real-time data analysis for data validation. Therefore, both NP and HEP have a common need to apply high-performance networks and near-detector HPC resources to obtain rapid feedback on event structure and detector systems performance. NP and HEP also share the Geant4 toolkit for detector simulation and the ROOT toolkit for statistical analysis; adaptation of these tools for efficient use of HPC resources would have a significant impact on both communities.

3.2.3.2 ASCR

Portable Software Stacks and Advanced Scheduling Mechanisms

Going forward, seamless access to compute resources is a cross-cutting requirement for many experimental sciences. While the growth in disk and archival storage requirements is relatively constant, the pattern for computing can be episodic. The local compute resources are scaled for an average need, to allow priority to shift between experiments as needed. However, often the campaign required to process data or generate simulations in a short time exceeds all local resources. For this reason, many of the experiments with large computational needs have made good use of distributed resources such as the OSG or (less often) commercial cloud resources. Given the nature of the experimental workloads, which can often fit into subsections of the large HPC resources concurrent with large-scale theory computations, a premium should be placed on portable software stacks and schedulers that can manage these mixed workloads.

Petascale at the Beam Line

Deploying local resources at the beam line and for local computing facilities that are smaller-scale versions of the large HPC resources will offer the opportunity, at the application level, to blur the distinction between online and offline resources. This will require some detailed understanding of the needs to balance I/O, compute, memory, and disk resources. A caution with the use case approach is that systems deployed locally are generally multi-use and may or may not be optimized for the workloads highlighted by the use case. Scaling from current use generally results in projections that are underestimated, reflecting the biases in the current budget-constrained systems. User applications are designed to adapt the resources that are available, and hardware is purchased to support existing applications—a barrier to innovation. One cross-cutting project is the instrumentation of systems to capture usage and bottlenecks to understand a more optimal ratio of I/O/compute/disk for future systems as one way of developing a data-driven dialog between application developers and hardware architects.

Emerging Needs in Data Science

There are emerging scientific questions that will require advanced machine learning techniques, possibly over combined data sets from multiple experiments or observations. In NP, an example of this is 3-D nucleon tomography, an emerging research area that will begin to acquire significant data samples with the JLab 12-giga electron volt (GeV) program (see Figure 3-13) and other experiments coming online, as well as integrating older data. The core of this program will be

global fits of worldwide data to theories with five or more parameters. Collecting and preparing the data for analysis and maintaining them for use over several decades will be required, as will provisioning for work flows that capture the parameters of the fits, the version of the theory codes, and any phenomenological preparation of the experimental data. Other disciplines will face similar challenges as theoretical and experimental advances open novel lines of inquiry. Appropriately capturing the data is one of the more pressing issues, because knowledge can be irretrievably lost. The NP community understands that neither ASCR nor NSF is currently positioned to provide large data repositories for the individual program offices. These program offices are likely to have to integrate and federate across their facilities' storage and/or integrate with the community repositories to meet the long-term scientific goals. There is a need to establish best practices and reference architectures (similar to the ESNNet DMZ architecture) as a way to guide the creation and maintenance of cost-effective and usable data repositories.



Figure 3-13. JLab aerial view after the \$338M 12-GeV upgrade showing the expanded experimental facilities. All experimental halls will be commissioned by May 2017 and will begin to take data at full rates. (Image used with permission from Thomas Jefferson National Accelerator Facility.)

Workforce

All domain sciences are increasingly data- and compute-driven and, for that reason, there is a need for user training and education in advanced computational techniques. In experimental HEP and NP, because of the timing of major facilities coming online in the early 1990s and limited career positions, we are experiencing the aging and retirement of a generation of laboratory-based physicists who specialized in computing for experimental physics. Their retirement presents an opportunity to refresh NP and HEP expertise and perhaps reconsider methods of collaboration and communication across the program offices, with ASCR playing a central role in providing expertise and best practices. A closely related concern is that of usability. Traditional “labware” often has a high barrier for usability, which can disenfranchise many physicists.

3.2.4 Computing and Data Needs

The computing needs and requirements of the experimental community are qualitatively different from those of flagship theoretical calculations that drive the massively parallel architectures of cutting-edge computing systems. The most significant difference is that experimental computing is dominated by the processing of the recorded detector signals, so that I/O becomes more of a focal point than sustained flops. This aspect extends beyond the data itself to the simulation of the experiments, which usually aims to emulate the recorded detector signals themselves, so that the same processing chain can be applied to both recorded and simulated data. As a result, experimental computing almost always involves the reading and/or the writing of large amounts of data. Managing those data sets and their complex processing work flows is a significant challenge. These considerations drive the computing needs and requirements for experiments.

The computing needs of experiments can be divided into several stages: onboard, online, semi-online, and offline. The onboard, online, and semi-online stages of analysis occur at the experimental site and are typically covered under project-specific resources. Here, only the offline computational needs of experiments, which can leverage centralized exascale computing facilities, are considered.

Figure 3-14 summarizes the computing resources required to accomplish the planned NP experimental programs. While the computing resources required to accomplish the theoretical objective of the NP program will yield an exascale computing ecosystem that is more than sufficient to accommodate the experimental computational requirements, experimental requirements are major drivers of the I/O design of an exascale system.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025
EXPERIMENT AND DATA




Figure 3-14. The capability and capacity computing resource requirements (left panel) and the hot and cold data requirements (right panel) in 2025 to accomplish the objectives of the NP Experiment and Data program.

3.3 Nuclear Structure and Reactions

3.3.1 Scientific Challenges and Opportunities

The field of nuclear structure and reactions (NS&R) research is now undergoing a renaissance, as a result of experimental developments and exciting theoretical progress in the low-energy nuclear many-body problem. The general challenge for this interdisciplinary field is to understand the principles of building up nuclear complexity out of fundamental degrees of freedom, which — when inspected at higher resolution — have a complicated structure of their own. The overarching questions that drive NS&R are as follows:

- Where do nuclei and elements originate?
- How are nuclei organized?
- How can nuclei be exploited to reveal the fundamental symmetries of nature?
- What are the practical and scientific uses of nuclei?

Complete answers to these questions require a much deeper understanding of atomic nuclei than is currently available. Both low-energy experimentation and nuclear theory address these questions in a synergic manner.

The planned NS&R research program at the exascale is perfectly aligned with the national nuclear physics program described in the *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) and the NRC decadal study (National Academies Press 2013) reports. A new generation of large-scale experiments and facilities using exotic nuclei to probe nuclear and particle physics is being constructed in the United States. FRIB is designed to study the structure of nuclei important for the processes that create all the elements in the universe. Nuclear theory at the exascale will calculate the properties and decays of atomic nuclei to understand the mechanism of nuclear binding and the formation of the elements and to extend these results to astrophysical environments that cannot be created in the laboratory. It will also address nuclear physics programs at JLab, which will delineate nuclear properties at short inter-nucleon distances and provide data on neutron distributions in nuclei.

$0\nu\beta\beta$ -decay experiments Majorana and nEXO, using ton-scale germanium and xenon detectors, are a high priority in the *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015; National Science Advisory Committee [NSAC] 2014[a]) and are being designed to (1) determine whether neutrinos are their own anti-particle (Majorana neutrinos) and (2) probe the neutrino's absolute mass scale in the laboratory. The envisioned exascale NS&R program will calculate and quantify uncertainties in nuclear $0\nu\beta\beta$ -decay matrix elements to guide next-generation $0\nu\beta\beta$ -decay experiments.

The Deep Underground Neutrino Experiment (DUNE) is being designed to accurately measure neutrino oscillations and the neutrino mass hierarchy (Fermilab undated[a], [b]). Here, nuclear calculations will describe neutrino and anti-neutrino scattering from argon to provide a more accurate determination of neutrino oscillation parameters, including the parity and CP violating phase.

Maximum success for these several hundred million- to billion-dollar experimental investments requires substantial simultaneous investment in computations of the structure and reactions of nuclei. As multiple DOE and NRC reports recognize, exascale computations are required to fulfill the critical role of guiding and making predictions for new experiments, and important computational advances are required to solve challenges expected at the exascale.

Figure 3-15 shows the degrees of freedom of NS&R research in the context of the theoretical roadmap of the nuclear many-body problem. QCD — the underlying theory of strong interactions — governs the dynamics and properties of quarks and gluons that form baryons and mesons; hence, QCD is also responsible for the complex inter-nucleon forces that bind nuclei. In this area, significant progress is being made by computing properties of the lightest nuclei and nuclear interactions within lattice QCD. A powerful suite of ab initio approaches based on inter-nucleon interactions provides a quantitative description of light- and medium-mass nuclei and their reactions. For medium-mass systems, global configuration-interaction (CI) methods that employ microscopic effective interactions offer detailed descriptions of nuclear excitations and decays. For heavy complex nuclei, the tool of choice is the nuclear DFT and its extensions.




Figure 3-15. Left: Case studies of the NS&R exascale program. Right: The basic elements that atomic nuclei are made of depend on the energy of the experimental probe and the distance scale. The building blocks of QCD are quarks and gluons, which form the mesons and baryons. In low-energy nuclear physics experiments, nuclei can be well described in terms of individual protons and neutrons, their densities and currents, or collective coordinates describing rotations and vibrations of the nucleus as a whole. Major theoretical approaches to the nuclear many-body problem (lattice QCD, ab initio methods, configuration interaction techniques, nuclear DFT, and collective models) are marked at different resolution scales. (Image reproduced with permission from Witke Nazarewicz; available in Bertsch et al. 2007.)

The NS&R exascale program will use different methods to perform validation and verification on nuclei in the regions of models' overlap to ensure accurate coverage of the entire nuclear landscape. The combination of these methods provides a coherent picture of the structure and reactions of all atomic nuclei over a wide range of energies and momenta. Each method has its major advantages: they can first be compared and validated in specific problems, and then each can be employed in the most appropriate regime. Although the methods described here are similar to those used for other strongly correlated quantum many-body systems, they are unique in that nuclear codes must solve for interactions that strongly couple different spin (up and down) and isospin (neutron and proton) states. NS&R computations involve extreme-scale codes that face common issues in efficient and reliable task-based parallelism, hierarchical computation and associated memory usage, and important issues in optimization at the largest scales. The community has successfully met such challenges over the past several generations of large-scale computing facilities through the SciDAC (DOE 2015) and INCITE programs (DOE Leadership Computing 2016).

What can modern nuclear structure theory do to improve its efficiency, enabling it to respond nimbly to opportunities for scientific discovery? One mechanism is to team up in large, multi-institutional efforts involving strong coupling among nuclear physics theory, computer science,

and applied mathematics. While experimentalists are no strangers to large collaborative efforts, nuclear theorists have eventually come to the realization that “the whole is greater than the sum of its parts.”

In this context, one representative example is the NUCLEI SciDAC collaboration (Ohio State University undated) involving nuclear structure theorists, computer scientists, and mathematicians from 16 institutions, including 10 universities and 6 national laboratories. The scope of nuclear science and math/computer science is quite wide-ranging. Over the course of NUCLEI and its Universal Nuclear Energy Density Functional (UNEDF) SciDAC predecessor, collaborations across domains have grown and now involve many direct connections. Figure 3-16 illustrates the present status of collaborations within NUCLEI. As is apparent from this network diagram, the math/computer science participants are directly embedded in the various nuclear theory efforts. In each partnership, these participants collaborate with physicists to remove barriers to progress on the computational/algorithmic physics side. This partnership has proven to be very successful, resulting in many excellent outcomes, some initially unanticipated. The UNEDF/NUCLEI case emphasizes the importance of assembling agile theory teams working on important questions and programmatic deliverables that would be difficult or impossible to tackle by individual investigators or atomized small groups.




Figure 3-16. Collaboration links within NUCLEI SciDAC, illustrating the integration of the project as of 2016. Physics and mathematics/computer science principal investigators are marked by blue and red dots, respectively. The boxes indicate areas of collaboration. Close NP/ASCR partnerships in SciDAC have been enormously successful and beneficial to the science goals. (Image produced by and used with permission from Randy Lusk.)


As in other areas of science, NS&R research uses a cycle of “observation-theory-prediction-experiment” to investigate phenomena, build knowledge, and define future investigations. Such an approach guides the relationship between theory and experiment: theory is modified or rejected based on new experimental data and the improved theory can be used to make predictions that guide future measurements. The positive feedback in the experiment-theory cycle, illustrated in a schematic way in Figure 3-17, can be enhanced if statistical methods and computational methodologies are applied to determine the uncertainties of model parameters and calculated observables. In partnership with applied mathematics and computer science, modern nuclear structure theory strives to estimate errors on predictions and assess extrapolations. This is essential

for developing predictive capability, as theoretical models are often applied to entirely new nuclear systems and conditions that are not accessible to experiment. Statistical tools can be used to both improve and eliminate a model or better define the range of a model's validity.

3.3.2 Priority Research Directions

The overarching goal of this research is to solve the quantum nuclear many-body problem — a quantum system of strongly interacting neutrons and protons — for both structure and important reactions. Because the applications and requirements are so broad, to solve this problem in its totality, the scientific and computational challenges will be addressed with different frameworks that are appropriate for different nuclear sizes and resolution, and for answering different questions in each regime (Figure 3-17). To this end, five priority research directions in exascale NS&R research have been identified and are outlined in the following paragraphs:

- Nuclei from the first principles,
- Light-nucleon reactions,
- Electroweak phenomena,
- Quantified heavy nuclei, and
- Dense nucleonic matter.



3.3.2.1 Nuclei from the First Principles

How Are Nuclei Organized?

In recent years, we have been witnessing the revolutionary development of complementary ab initio methods based on realistic nuclear forces tuned to reproduce selected scattering and bound-state properties, and operators that are consistent with interactions employed. First-principles computations have been applied to critical problems in light nuclei, including the structure of the Hoyle state in ^{12}C , which makes possible the nucleosynthesis of the light elements that enable life; the long half-life of ^{14}C , which enables carbon dating; and the interaction between neutrinos and nuclei. Ab initio theory is now reaching medium-mass nuclei such as ^{52}Ca , ^{78}Ni , and even heavier semi-magic systems (Figure 3-18). The results of quantified first-principles calculations will yield nuclear properties that will be directly confronted with experimental data on exotic nuclei from FRIB and short-range correlation and parity-violating electron scattering experiments at JLab.

The target problem for the next decade will be to study the physics of neutron-rich nuclei that are important for the creation of all the heavy elements in the universe.




Figure 3-18. Nuclei calculated in ab initio approaches prior to 2005 (upper panel) and now (lower panel). Exascale computing will enable ab initio calculations for many of the nuclei critical to the synthesis of the heavy elements. (Image available in Hergert et al. 2016.)

Path to Exascale in Nuclei from the First Principles

- 2016: Obtain solutions of light- and medium-mass nuclei with realistic nucleon–nucleon and three-nucleon interactions.
- 2020: Predict properties of nuclei far from stability relevant to FRIB experiments and their electromagnetic transitions using operators consistent with the strong interactions.
- 2025: Predict properties of heavy nuclei and investigate the shell structure of nuclei heavier than ^{208}Pb .

3.3.2.2 Light-Nucleon Reactions

How Did Visible Matter Come into Being and How Does It Evolve?

One of the outstanding challenges in contemporary nuclear theory is to describe nuclear structure, clustering, and reactions within a single coherent framework. Besides its intrinsic intellectual merit relating to our understanding of the strong and electroweak interactions in nuclei, accurate predictions are needed for difficult-to-measure reactions that synthesize the elements in stars or come into play in terrestrial applications of nuclear fusion. Ab initio approaches have made great strides toward meeting this challenge in lighter nuclei by enabling calculation of a variety of strong and electroweak processes in a realistic framework. Examples are key reactions involving composite projectiles, such as



or



Further developments and exascale resources are critical to improving the accuracy and extending the range of applicability of these methods to heavier systems and more complex reaction processes (Figure 3-19), so that ‘measurements’ from simulations can be combined with experiments to solve outstanding problems in nuclear astrophysics. The target problem for the next decade will be to arrive at a quantitative description of binary and ternary reaction processes in systems with up to $A = 20$ nucleons and beyond; among these are the $2\alpha(\alpha,\gamma){}^{12}\text{C}$ and ${}^{12}\text{C}(\alpha,\gamma){}^{16}\text{O}$ reactions critical for the formation of life and the synthesis of heavier elements in stars.




Figure 3-19. Capture of ^3He by alpha particles to form ^7Be . Ab initio techniques make it possible to calculate low-energy reactions, and exascale computing resources will enable the calculation of many important processes that are difficult to measure in the laboratory. (Image reproduced with permission from Sofia Quaglioni; available in Doherty et al. 2016.)

Path to Exascale in Light-Nucleon Reactions

- 2016: Make predictions of solar pp-chain reactions with realistic nucleon-nucleon and three-nucleon interactions.
- 2020: Make quantified predictions of nucleon-capture reactions (e.g., in the carbon-nitrogen-oxygen cycle) up to mass $A = 20$.
- 2025: Make quantified prediction of the α -capture reactions that synthesize carbon, oxygen, and heavier elements in stars.

3.3.2.3 Electroweak Phenomena

How Can Nuclei Be Exploited to Reveal the Fundamental Symmetries of Nature?

The weak and strong interactions are inextricably linked, and a thorough understanding of the many-body physics underlying nuclear structure and dynamics is required to extract neutrino physics and address the role of weak interactions in astrophysics and fundamental symmetries. The NS&R exascale program will address electroweak phenomena in a large set of nuclei over a broad range of energies. At low energies, this research will (1) provide accurate quantitative estimates of nuclear matrix elements that are critical for the interpretation of $0\nu\beta\beta$ -decay experiments Majorana and nEXO, (2) enhance our knowledge of weak transitions relevant to tests of physics beyond the Standard Model, and (3) allow us to quantitatively understand neutrino propagation in the dense matter in CCSNe and NSs. At higher energies, exascale capabilities will provide essential input for unraveling the response of neutrino detectors in experiments, such as DUNE—a high priority in the DOE P5 report (DOE 2014)—which will measure neutrino oscillation parameters, including the CP violating phase in the neutrino mass mixing matrix that makes neutrinos and antineutrinos oscillate differently, and that will determine the neutrino mass hierarchy (Fermilab undated[a], [b]). The target problems for the next decade will be to (1) carry out ab initio calculations of weak transitions between low-lying states of light- and medium-mass nuclei; (2) develop response functions describing inclusive low- (few to tens of MeV) and high-energy (hundreds of MeV to a GeV) neutrino scattering induced by neutral and charge-changing weak currents from nuclear targets (such as ^{12}C and ^{40}Ar) and low-energy neutrino propagation in nucleonic matter; and

(3) examine the matrix elements for $0\nu\beta\beta$ decay of ^{76}Ge , ^{136}Xe , and other nuclei will be used in $0\nu\beta\beta$ experiments (Figure 3-20).




Figure 3-20. Success and challenges in electroweak nuclear physics. Electron scattering from nuclei can be calculated using ab initio techniques (left panel), but the uncertainties in $0\nu\beta\beta$ -decay matrix elements remain a challenge (right panel). Exascale computing resources will be used to dramatically reduce uncertainties in these predicted rates. (Image on left available in Lovato et al. 2016. Image on right produced by and used with permission from Jangling Yao.)

Path to Exascale in Electroweak Phenomena

- **2016:** Determine the ab initio electroweak response of ^{12}C , with quantified uncertainty.
- **2020:** Identify DFT-based and ab initio matrix elements for $0\nu\beta\beta$ decay of several nuclei with many-body currents and quantified uncertainty.
- **2025:** Determine ab initio electroweak response of ^{40}Ar and identify ab initio matrix elements for $0\nu\beta\beta$ decay of ^{76}Ge , ^{136}Xe , and ^{130}Te with uncertainties of less than 25%.

3.3.2.4 Quantified Heavy Nuclei

How Are Nuclei Made and Organized? How Are Heavy Elements Produced in Stars?

The application of HPC has revolutionized theories of complex nuclei by both optimizing the input and carrying out advanced applications. Here the tool of choice is DFT, which can be benchmarked with ab initio computations in the calcium region. The time-independent and time-dependent DFT extensions have provided quantitative descriptions of one of the toughest problems of nuclear structure: nuclear large-amplitude collective motion, which includes such phenomena as shape coexistence, fission, and heavy-ion fusion.

Recent highlights include large-scale surveys of global nuclear properties that enabled predictions of nucleonic drip lines with quantified uncertainties, the elucidation of the structure of superheavy nuclei, quantitative modeling of fission and heavy-ion fusion reactions, and uncertainty quantification of nuclear data for astrophysical r-processes. As of today, the typical deviations from experiments are of the order of 700 keV for nuclear masses, about 2 orders of magnitude for spontaneous fission half-lives (over a range of more than 35 orders of magnitude), and approximately 30% for fission mass distributions in actinide nuclei. The scientific goal for the next decade is to bring these values down to 300 keV for masses, a factor of 2 for fission half-lives, and 10% or less for yield distributions (Figure 3-21). This will require the development of more realistic

energy functionals, which will include finite-range and/or three-body nuclear forces. The target problem for the next decade will be to develop quantum mechanical tools for precision nuclear spectroscopy and determine the ground state and decay properties of all atomic nuclei involved in nucleosynthesis.




Figure 3-21. Spontaneous fission yields of ^{240}Pu predicted in nuclear DFT. Exascale computing will enable more accurate treatments of spontaneous and induced fission, which are important in many applications and in r-process nucleosynthesis. (Image used with permission from Witek Nazarewicz; available in Sadhukhan et al. 2016.)

Path to Exascale in Quantified Heavy Nuclei

- 2016: Determine quantified masses, spontaneous fission half-lives, and yields for even-even nuclei.
- 2020: Determine neutron-induced fission and $0\nu\beta\beta$ rates with microscopically constrained nuclear energy functionals.
- 2025: Identify global nuclear physics inputs, with theoretical uncertainties, for r-process nuclei.

3.3.2.5 Dense Nucleonic Matter

What is the Nature of Dense Nucleonic Matter in the Cosmos and on Earth?

The properties of dense matter are critical to the study of compact objects in nuclear astrophysics, including especially NSs and CCSNe. These environments form the key to creating all the heavy elements in the universe and probing matter at extremely high densities and low temperatures; they are the keys to the phase diagram of QCD in this regime. To understand these astrophysical objects, one needs to know the EoS of dense matter and also the response to electroweak probes such as neutrinos. To this end, an interdisciplinary approach is essential to integrating low-energy nuclear experiments and theory with knowledge from astrophysics, atomic physics, computational science, and electromagnetic and gravitational-wave astronomy. The input of the exascale NS&R program to this mix is essential. It includes ab initio and DFT approaches to the EoS of nuclear matter, studies of nuclear matter at both supranuclear and subnuclear densities, as well as development of electroweak current models to provide a realistic description of the dynamic properties of dense matter (Figure 3-22). The target problems for the next decade will be to determine the EoS for cold dense matter for NSs and supernovae; calculate neutrino interactions in dense matter for supernovae spectra and proto-NS cooling; and understand superfluidity in dense matter for NS cooling.




Figure 3-22. Mass-radius relation for NSs. Exascale computing resources will enable more accurate studies of dense matter, including the EoS and neutrino opacities important in NSs and supernovae. (Image used with permission from Joe Carlson; available in Galdolfi et al. 2012.)

Path to Exascale in Dense Nucleonic Matter

- 2016: Calculate the zero temperature EoS of neutron matter.
- 2020: Explore superfluid pairing gaps in dense matter and EoS of nuclear matter at finite temperature.
- 2025: Determine the neutrino response of nuclear matter at zero and finite temperature.

3.3.3 Cross-Cutting Issues and Opportunities

The physics of nuclei is closely related to studies at several major experimental facilities under construction or being planned in the United States. FRIB is focusing on NS&R, but also has nuclear astrophysics as a major focus. Exascale computing resources in NS&R include many-body theory of dense nucleonic matter as found in NSs and CCSNe. The EoS and the propagation of neutrinos in dense matter are critical components of understanding these astrophysical objects. In addition, the structure of the NS crust and reaction rates there, and more generally, are a major overlap with nuclear astrophysics.


At the smallest length scales, NS&R has very important overlaps with lattice QCD and with BSM physics. Lattice QCD can potentially refine the nuclear interactions and currents beyond experiment alone—for example, the three-neutron interaction. Lattice QCD can also provide valuable information about multi-nucleon electroweak currents. NS&R plays an important role in studies of neutrinos and BSM physics. $0\nu\beta\beta$ -decay experiments are searching for a new lepton-number violating process, the observation of which would identify the Majorana nature of neutrinos. There may also be other routes to this new physics. Accelerator neutrino oscillation experiments rely on an understanding of neutrino-nucleus cross-sections and their energy dependence. Two-nucleon correlations and currents play a critical role in these processes; hence, a sophisticated understanding is required to extract the CP-violating phase in the lepton sector and the neutrino mass hierarchy.

Finally, nuclear theory has many areas of overlap with studies of strongly correlated quantum matter in many disciplines, including cold atoms, condensed matter, and electronic structure. The methods and algorithms used in these studies, although different in implementation, all have analogues in these other fields. These analogous methods imply that software and architectural advances valuable to the nuclear structure community can also have wide applicability. Examples include sparse and dense linear algebra, task-based parallel algorithms and scalability, hierarchical

memory management, and optimization and uncertainty quantification. There are also similarities in the physics being addressed, including superfluidity in homogeneous and inhomogeneous systems, transitions from ab initio to DFT, transport coefficients in complex systems, and many more.

As shown in Figure 3-23, the atomic nucleus is placed at the center of the Quantum Ladder: it provides a connection between the smallest and the largest. Being a congregation of neutrons and protons, it emerges from complex interactions between quarks and gluons on a scale of femtometers (fm). But its properties also determine the behavior of giant stars on a gigameter scale. Indeed, nuclear structure encompasses phenomena over an incredibly wide range of energies and distances. Atomic nuclei are laboratories of fundamental laws of nature and thus are linked to particle physics. At the same time, nuclei exhibit behaviors that are emergent in nature and present in other complex systems studied by condensed matter physicists and quantum chemists. Cross-cutting challenges of NS&R exascale research include the fermionic sign problem, treatment of finite-size effects, load balancing, and many-body approaches to open quantum systems (e.g., quantum Monte Carlo [QMC], CI, coupled clusters [CC], DFT).

The exascale NS&R research program will have high societal relevance. The coupling of NS&R with HPC will provide unique opportunities for training of future science leaders, new directions in education, and outreach to the public. Precise calculations on fusion of light nuclei and nuclear fission will have impact on energy research, as well as defense, waste transmutation, and stockpile stewardship programs. Predicting properties of specific nuclei, with characteristics adjusted to specific research needs, will have an impact on medical and biological research, materials science, environmental science, and other areas. With a picture of nuclei based on the correct microphysics, exascale computing, and experiment, we will develop a comprehensive model of the atomic nucleus: to understand, predict, and use.




3.3.4 Computing Needs and Requirements

The NS&R community has recognized several important requirements for an exascale ecosystem. These are presented in detail in other sections of this report. A major need across all subareas in NS&R is continued access to the latest-generation machines. The need for increased core and node hours is dramatic. Certain specific applications, including CI and reactions based upon that, also require large aggregate memory. Another major element is a strong and diverse workforce, including physicists, mathematicians, and computer scientists, as well as domain scientists with a

background in these areas. Significant effort is required to sustain a system of this type, particularly in providing an attractive and permanent career path for scientists with a background in both physics and computing. Many other important needs have been identified, including I/O and storage needs. Figure 3-24 summarizes the computing resources required to accomplish the planned NS&R program.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025
NS&R



* Exaflop-system-year refers to the total amount of computation produced by an exascale computer in 1 year.

Figure 3-24. The capability and capacity computing resource requirements (left panel) and the hot and cold data requirements (right panel) in 2025 to accomplish the science objectives of the NS&R program.

3.4 Cold Quantum Chromodynamics

3.4.1 Scientific Challenges and Opportunities




Figure 3-25. Cover of *Nature* magazine containing an article that discusses the conflicting experimental measurements of the proton radius (Nature 2010).

A precise understanding of the proton, the neutron, and the forces between them is critical to many aspects of subatomic science — from the earliest moments of our universe and the fundamental conditions necessary for life, to national security, the production of carbon-free energy, and the transmutation of nuclear waste. In each of these areas, and beyond, there are clear needs to improve our present understanding of, and to enhance our predictive capabilities in, nuclear systems — from the very smallest nucleus and its reactions to nuclear matter found in the dense interiors of CCSNe. In designing the next generation of HEP proton colliders, the particle production rates and their discovery potential depend sensitively on the structure of the proton. The present uncertainties in the structure of the proton translate into uncertainties in machine and detector design parameters and projected experimental costs.

In designing experiments to search for the violation of lepton number (L) — that may manifest itself, for instance, in the $0\nu\beta\beta$ decay of nuclei — uncertainties in the expected nuclear decay rates currently result in

significant uncertainties in estimates of the amount of the active material required to accomplish the experimental objectives. These uncertainties have implications for precision atomic physics. The electronic structure of atoms is sensitive to the distribution of charge within the nucleus at its center, and recent experimental results have demonstrated that the radius of the proton is not as well determined as scientists had once believed. The approximately 7σ discrepancy between electronic and muonic measurements is under active investigation (see, for example, Figure 3-25), and the computational technology now exists to calculate the proton radius from first principles, complementing the experimental program. The present-day uncertainties in nuclear reactions and the behavior of matter under extreme conditions continue to be addressed by ever-increasingly precise determinations of the spectra of hadrons, their structure, and their interactions. Quantities that are the foci of investigations in cold QCD — the structure and reactions of mesons, nucleons, and light nuclei — have an impact on an array of systems, from uncertainties in fusion reactions that play an important role in the operation of the NIF (Lawrence Livermore National Laboratory undated); to uncertainties in isotope and neutrino production in nuclear fission reactors; to neutrino interaction rates in experiments, such as DUNE (Fermilab 2016), designed to determine the fundamental properties of neutrinos; to the behavior of matter in explosive astrophysical environments; to expected features of the gravitational waveforms from the inspiral of binary NSs that are expected to be seen using LIGO (Figure 3-26). Results in light nuclei, combined with efforts of many-body theorists, are expected to have a direct impact on the studies of many-nucleon systems, including those studied at the FRIB and through astrophysical observations and those used in accelerator neutrino and $0\nu\beta\beta$ -decay experiments. The anticipated advances in high-performance computing and data ecosystems (HPCDEs) that are expected in the exascale era are required to provide precise calculations of key quantities and systems necessary to accomplish the scientific objectives of the cold QCD program.



The drive toward a precise theoretical understanding of cold QCD systems has traditionally spurred the development of HPC. It was the desire to perform increasingly more realistic QCD calculations in HEP that led to the development and fabrication of the QCD on digital signal processors (QCDSP) then QCD on a chip (QCDOC) by the Lattice QCD Group at Columbia University, in collaboration with the University of Edinburgh and BNL, and ultimately to IBM's Blue Gene/L supercomputer. This trend continued with the deployment of large, GPU-based supercomputers, such as the one deployed by the USQCD consortium at JLab in 2010 and dedicated to lattice QCD calculations.

Protons, neutrons, nuclei, and all strongly interacting particles (hadrons) emerge from the dynamics and interactions of quarks and gluons, fundamental building blocks of our universe, dictated by QCD. Hadrons are composite systems comprising “valence” quarks, immersed in a “sea” of quarks, antiquarks, and gluons that are quantum mechanically fluctuating in and out of the vacuum. They are entangled quantum systems with an indefinite particle number, and the intrinsically nonlinear and quantum mechanical nature of QCD confines the quarks and gluons into hadrons so that they have never been observed in isolation. In analogy with electromagnetism, gluons carry the force between quarks in a way that is similar to photons carrying the force between electrically charged particles. The nonlinearity in QCD arises from the gluons carrying “color” charges and therefore interacting with themselves, with significant consequences. For instance, the strength of interaction between quarks becomes stronger at larger distance, to the point where the vacuum will “spark” to produce quark–antiquark pairs when their separation becomes on the order of 1 fm.

Further, symmetries of the underlying laws responsible for the dynamics of the quarks and gluons are masked by the structure of the strong vacuum, and hadrons and nuclei are emergent phenomena resulting from the simple laws obeyed by the quarks and gluons. Reproducing and predicting the spectra, structure, and interactions of the hadrons, and the nuclei that are bound states of protons and neutrons, without the use of HPC has so far eluded scientists. The impact of the nonlinearities is greatly enhanced by quantum fluctuations, rendering analytic techniques of limited utility for

quantitatively describing low-energy processes. Soon after the discovery of QCD as a candidate theory describing strong interaction processes, Wilson (1974) formulated lattice QCD as a numerical method to provide the solution to QCD in the low-energy regime.



Figure 3-27. Cover of the NSAC 2015 Long-Range Plan for Nuclear Science.

The United States operates world-leading and complementary accelerator facilities for research into cold QCD. The CEBAF at JLab is the world's most powerful facility of its kind for studying cold QCD systems, delivering multi-GeV electron beams onto nuclear targets for a precision probe of “cold” nuclear matter (JLab undated[a]). This facility recently underwent an upgrade, at a cost exceeding \$300 million, to increase its beam energy from 6 GeV to 12 GeV. The facility enables precision investigations of the valence structure of the nucleon and nuclei and of the spectra of mesons and baryons and promises to uncover and characterize new exotic states of matter. The \$730-million FRIB

is currently under construction at MSU and is scheduled to become operational in 2022 (MSU undated). The experimental program that is planned for FRIB will probe the limits of stability of neutron-rich nuclei and allow for a refinement of the nuclear forces that dictate the nature of high-density matter, which is present in the interior of NSs and which imprints itself onto the gravity-wave signals from the inspiraling NSNS systems that are expected to be observed using LIGO (California Institute of Technology undated). The NSAC (DOE SC undated) *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) recommends the evolution of the JLab and RHIC (BNL undated) programs into one centered around an EIC (Cartlidge 2015), with a focus on unraveling the gluonic and “sea” structure of nucleons and nuclei (Figure 3-27). This plan also recommended the construction of a ton-scale $0\nu\beta\beta$ -decay detector to search for lepton number violation manifesting itself through the $0\nu\beta\beta$ decay of certain nuclei.

Precision lattice QCD studies of the structure, spectra, and interactions of mesons, baryons, and nuclei, and of the fundamental symmetries of nature, are now crucial to the U.S. nuclear physics research program and also to components of the HEP research program. This field is entering an era in which lattice QCD calculations will be routinely performed at the physical quark masses and include electromagnetism. The program of computation is closely tied to the U.S. experimental program in nuclear physics and to the scientific milestones established by the NSAC and provides crucial support to other areas of nuclear and particle theory. The component of the program focused on the spectrum and structure of the mesons and baryons, including the study of exotic states, is crucial in optimizing the scientific productivity of the 12-GeV experimental program at JLab. It is evolving to include anticipated needs of a future EIC. The program on light nuclei, hypernuclei, and nuclear forces complements and supports the FRIB facility and provides input for studies, combined with nuclear many-body theory, of dense astrophysical environments. The main focus is to refine the chiral nuclear forces, including multi-neutron forces, defined through low-energy effective field theories (EFT), which are used as input into nuclear many-body calculations of nuclear structure and reactions. The fundamental symmetries component — which is intimately connected to the hadron structure and nuclear forces program — is essential to understanding the origin of time-reversal (T) violation once a neutron electric dipole moment (nEDM) is observed and to searches for nonstandard model physics at neutron facilities at ORNL (ORNL undated) and Los Alamos National Laboratory (LANL undated). It is also important to improving estimates of L-violating $0\nu\beta\beta$ -decay rates of nuclei in support of the planned ton-scale $0\nu\beta\beta$ -decay experiment.

3.4.2 Priority Research Directions

During the next decade, as cold QCD moves through the pre-exascale era into the exascale era, we expect nuclear physics to be revolutionized by the newfound ability to compute low-energy strong interaction processes with high precision that can be directly compared with experiment. We fully anticipate that, during this period, there will be important quantities for which the lattice QCD calculations supersede experimental measurements in precision and that are significantly less expensive to perform. The *2015 Long Range Plan for Nuclear Science* identifies the study of the gluonic structure of the nucleon and nuclei through the construction of an EIC as an emerging priority in QCD research (DOE and NSF 2015). The evolution of the JLab and Brookhaven National Laboratory QCD programs into an EIC program, along with the significant growth of the fundamental symmetries experimental program through the ton-scale $0\nu\beta\beta$ detector, the theoretical workforce that is currently focused on hot and cold QCD is expected to realign itself to optimally support the EIC and fundamental symmetries experimental programs of the next decade. We have selected five representative priority research directions in the area of cold QCD:

- Hadron structure;
- Hadron spectroscopy;
- Nuclear physics from lattice QCD;
- Fundamental symmetries and new physics; and
- Simulating the quantum vacuum: the generation of gluon field configurations.

Each of these is discussed in the sections that follow.

3.4.2.1 Hadron Structure

A precise knowledge of the structure of the light hadrons, such as the pion, proton, and neutron, is essential for many aspects of subatomic physics, from the electronic structure of atoms, to the response of detectors to neutrinos, to the design of the next-generation hadron colliders. Such knowledge includes the complete and precise decomposition of the mass of the proton, the decomposition of its spin, and a complete 3-D tomography of the nucleon to be realized through the determination of generalized parton distribution functions (GPDs) and transverse momentum dependent distribution functions (TMDs). These studies support and complement the extensive experimental programs at JLab, at RHIC, and at CERN's LHC and will be essential in mapping the gluonic structure of nucleons and nuclei in the EIC era (a cartoon of which is shown in Figure 3-28). DOE is currently supporting the Coordinated Theoretical Approach to Transverse-Momentum Dependent Hadron Structure in QCD Topical Collaboration as part of its support for research into the structure of the nucleon (recent results obtained for the TMDs are shown in Figure 3-29). The nature and strength of the interactions of the light hadrons with the carriers of the weak force, dictated by their structure, is critical to the success of experiments focused on determining the properties of neutrinos. Related matrix elements play a central role in searches for dark matter and other physics BSM.




Figure 3-28. Lattice QCD calculations of the gluonic structure of the proton that are anticipated in the exascale era will support and complement an EIC experimental program envisioned as the future of QCD research in the United States. (Figure reproduced with permission from David Richards.)

During the pre-exascale era, hadron structure calculations will precisely determine fundamental quantities characterizing the nucleon, including form factors, moments of parton density, helicity, and transversity distributions, and moments of the GPDs and TMDs, providing a 3-D map of the nucleon. These calculations will directly support and complement experiments in JLab's 12-GeV program, at RHIC-spin, and at Fermilab. With the JLab and RHIC experimental programs evolving into an EIC program in the future, the theoretical research will increasingly emphasize the gluonic structure of the nucleon and nuclei and more generally the role that gluons play in matter. Theoretical calculations will play a crucial role in designing an EIC and in the subsequent analysis of experimental data that it produces. For technical reasons, such lattice QCD calculations are significantly more complex than those of the quark structure. A particularly important experimental conflict to be resolved is the discrepancy between measurements of the proton charge radius as measured with electrons and muons. Is this discrepancy revealing non-universal interactions of the leptons? This question remains to be answered and can, in fact, be resolved during the next several years with exascale computing resources.




Figure 3-29. Present-day calculations of the proton TMDs for the up and down quarks at $x=0.1$. k_x and k_y correspond to the transverse momenta. (Image available in Accardi et al. 2012.)

With exascale resources, calculations of the structure of hadrons will enter a high-precision era, essentially coinciding with the expected turn-on of the experimental program at an EIC. The valence-quark observables will be calculated with extremely high precision where required, such as g_A and its associated form factor. They will become calculable to better than one part in a thousand

— a precision required to refine calculations of the pp-fusion rate that dictates solar burning. These calculations will include fully dynamical quantum electrodynamics and strong isospin breaking. The challenges in calculating gluonic observables suggests that supporting and complementing the EIC program can be performed only during the time frame of operation of the EIC.

Path to Exascale in Hadron Structure

- **2016: Determine precise distributions of spin, charge, and currents in the nucleon and obtain an increasingly precise determination of the charge radius of the proton.**
- **2020: Complete detailed mapping of the 3-D structure of the nucleon and the first comprehensive exploration of its gluonic structure. Complete mapping of the quark and gluon energies confined within the proton.**
- **2025: Complete precision calculations of nucleon couplings and associated form factors to all quark and gluon interactions. Obtain a detailed understanding of the gluonic contributions to the mass and spin of the proton, nuclei, and exotic states.**

3.4.2.2 Hadron Spectroscopy

Hadron spectroscopy, the study of the spectrum of mesons and baryons, plays a pivotal role in developing our understanding of QCD. Lattice QCD calculations are critical to this effort, suggesting the existence of new types of particles, and address a unique challenge to our understanding of QCD: how does the complexity of the hadron spectrum emerge from the interactions of the quarks and gluons? In particular, a key goal in the study of cold QCD is to identify the presence of “hybrid” mesons and baryons, exotic states of QCD in which the gluonic degrees of freedom are manifest, and whose discovery is the primary motivation of the GlueX experiment at JLab (JLab undated[b]). Results from recent experiments have upended the conventional picture that hadrons have a simple classification scheme (for example, Shepherd et al. 2016), which has ignited significant interest.

Heading into the exascale era, a key objective of hadron spectroscopy is the determination of resonance properties of the hadron spectrum. The restriction imposed upon lattice QCD calculations by the finite spatial volumes of the space-time lattice is a feature that enables a mapping of the energy dependence of scattering amplitudes from which resonance properties can be determined. Predictions of the spectrum of mesons, including exotic states (schematically shown in Figure 3-30), their strong decays, and their photo-couplings obtained from these calculations will guide experimental searches and measurements at the GlueX experiment and confront future predictions in other sectors of QCD (a recent calculation is shown in Figure 3-31). Notably, in the charmonium sector of QCD, future calculations will address the origin of new and unexpected experimentally observed resonant structures that require a new understanding of the emergent dynamical degrees of freedom from the quarks and gluons. Calculations will address questions requiring high precision in these systems in which electromagnetism is important. The technical challenge facing this program is the computation of the highly excited, finite-volume energy spectrum directly from QCD. Highly scalable source construction techniques are needed for the anticipated calculations, as well as a new software and work-flow infrastructure that can support the intense I/O demands placed on the computing systems.




Figure 3-30. The structure of conventional and hybrid mesons. (Figure reproduced with permission from Robert Edwards and JLab.)

Scientists have found that the recently discovered plethora of unconventional states (Shepherd et al. 2016) are characterized by unnaturally small energy scales, similar to those encountered in the physics of nuclei. While remarkable progress is expected in deciphering their nature and identifying what is expected to be towers of partner states, high-precision calculations of these systems will require multi-exascale computing resources.




Figure 3-31. Isoscalar and isovector meson spectra at a pion mass of 391 MeV and spatial volume of 3 fm (Dudek et al. 2013). States enclosed in orange rectangles are the lowest lying states with intrinsic gluonic structure. (Figure reproduced with permission from Robert Edwards.)

Path to Exascale in Hadron Spectroscopy

- 2016: Search for exotic hadrons in QCD and study their resonance properties to guide the GlueX experiment at JLab.
- 2020: Complete the new hadron-classification scheme recently discovered with charmed quarks.
- 2025: Complete a quantitative exploration of the strong-interaction glue, and its excitations, that binds quarks into hadrons, to support and complement an EIC experimental program.

3.4.2.3 Nuclear Physics from Lattice QCD

Lattice QCD calculations play a critical role in improving nuclear forces and interactions. Providing the fundamental input into nuclear many-body calculations, these forces are essential in predicting

the structure and decays of nuclei, nuclear reaction rates, the response of nuclei to electroweak probes and candidate dark matter particles, the rates of $\beta\beta$ and $0\nu\beta\beta$ decay of nuclei, and the behavior of matter in extreme astrophysical environments. Some quantities that are difficult (or impossible) to access experimentally can be determined using lattice QCD calculations, thereby providing the theoretical support that is critical in capitalizing on the investments made in the U.S. experimental program. The forces that describe the stable and long-lived nuclei have been tightly constrained through decades of experimentation. However, the present uncertainties in the nuclear forces that determine the structure and behavior of neutron-rich nuclei, and of the matter that forms during the core-collapse of supernovae and other dense environments, are limiting the precision of predictions for the structure and dynamics of these systems. The FRIB experimental program is designed to reduce these uncertainties through the direct measurement of processes involving short-lived, neutron-rich nuclei. To complement this effort, lattice QCD calculations aim to refine other components of the forces through studies of neutron-rich systems and hypernuclei that are not accessible in experiment. Similarly, the present (and larger) uncertainties in the multi-nucleon interactions with electroweak currents and processes resulting from BSM physics, are limiting our ability to predict the response of nuclei to fluxes of neutrinos and their rate of decay in the presence of L-violating processes.

Entering the exascale era, it is the lightest nuclei and systems of several nucleons that will be studied, and the electroweak interactions in these systems are required to be known at the percent level. Further refinement of the nuclear forces, through extremely precise calculations of nucleon–nucleon (NN) interactions over a wide range of angular momenta, and also high-precision calculations of multi-neutron interactions, including electromagnetism, and precision electroweak studies will require multi-exascale resources. This level of precision is required for first-principles calculations of the structure and behavior of dense matter at the densities generated in CCSNe. Reproducing the nuclear interactions and scattering parameters from first principles will define a watershed moment in history for nuclear physics, forever changing the nature of this field. Significant progress has been made in calculating the properties of light nuclei with lattice QCD (Beane et al. 2013), with the binding energies of the s-shell nuclei and hypernuclei calculated at a number of values of the pion mass (the results of recent calculations are shown in Figure 3-32), along with nucleon–nucleon and nucleon–hyperon scattering phase shifts, and nuclear magnetic moments (recent results are shown in Figure 3-33) and polarizabilities (Beane et al. 2014). Recently, the first lattice QCD calculation of an inelastic nuclear reaction, $np \rightarrow d\gamma$, was accomplished (Beane et al. 2015). The most significant difficulty encountered in multi-nucleon calculations is the long-time behavior of the requisite correlation functions that exhibit an exponentially decaying signal-to-noise ratio resulting from general quantum mechanical principles. Improvements in the source and sink structures will be a focus of algorithm design in the pre-exascale era.




Figure 3-32. The spectrum of light nuclei and hypernuclei when the up-, down-, and strange-quarks have equal mass (Beane et al. 2013). (Image reproduced with permission from the NPLQCD Collaboration.)




Figure 3-33. The magnetic moments of the light nuclei at the SU(3) symmetric point (blue bands) and the corresponding experimental values (red dashed lines) (Beane et al. 2014). (Image reproduced with permission from the NPLQCD Collaboration.)

Through the exascale era, the properties and interactions of light nuclei, multi-nucleon, and neutron-rich systems, as well as hypernuclei, will be calculated at and near the physical quark masses and will include electromagnetism — leading to a significant improvement in the precision of our description of the nuclear forces. The anticipated lattice QCD calculations will enable a complete dissection of the chiral nuclear forces at the levels required for precision calculations in nuclei and dense matter. Through collaboration with nuclear many-body theorists, the results of these studies will be incorporated into nuclear many-body calculations of quantities of critical importance in

nuclear physics. Not only will they support and complement the experimental programs planned for JLab, FRIB (Figure 3-34), the EIC, LBNL/DUNE, and searches for the $0\nu\beta\beta$ of nuclei, but they will also lead to a reduction in uncertainties in rates of processes important for the generation of energy. In addition to refining the multi-body interactions between protons and neutrons, this extensive series of calculations will tightly constrain the three-neutron and four-neutron interactions and provide precision hyperon–nucleon forces — quantities that are vital in predicting the behavior of dense matter, but that have so far proven elusive in laboratory experiments.

Isotopes for Applications from FRIB Fragment Separator




Figure 3-34. A schematic of the isotope separation at FRIB. (Image produced and used with permission from the Experimental Systems Division, FRIB Project.)

The responses of nuclei to electroweak probes and candidate dark matter particles are key elements in the analysis of neutrino and dark matter direct detection experiments; the multi-body aspects of these interactions will be precisely determined, enabling a reduction of the uncertainties in processes such as neutrino–nucleus reactions.

Path to Exascale in Nuclear Physics from Lattice QCD

- **2016:** Determine two-nucleon and hyperon–nucleon scattering, three-neutron forces, and the electromagnetic structure of light nuclei at unphysical quark masses. Preliminary calculations at the physical quark masses are being performed in Japan and Europe.
- **2020:** Complete U.S. calculations at the physical quark masses. Determine constraints on chiral nuclear forces for input into nuclear many-body calculations.
- **2025:** Perform precision calculations of light nuclei and chiral nuclear forces at the physical quark masses, including electromagnetism. Complete precision studies of the gluonic contribution to the mass and structure of light nuclei.

3.4.2.4 Fundamental Symmetries and New Physics

Generating the observed asymmetry between matter and antimatter in the universe requires the non-conservation of both baryon number (B) and CP during non-equilibrium dynamics in the earliest moments of the universe. The CP and B+L violation present in the standard model is insufficient to produce the observed asymmetry, and a number of current and planned experiments are designed to explore T, CP, B, and L violating processes in an effort to identify and understand the additional sources of violation that must be present. In particular, DUNE, which is expected to start operation in 2024, will search for B-violating proton decay, and neutron–antineutron oscillation experiments are also envisaged. An experiment to measure the nEDM (a cartoon of which is shown in Figure 3-35), a probe of T violation, is under development at the Spallation Neutron Source (SNS) at ORNL (ORNL undated). A new flagship initiative in nuclear physics is the development and deployment of a ton-scale experiment to search for $0\nu\beta\beta$ of nuclei that, if observed, would provide explicit evidence of L violation. In order to constrain the nature of the associated new physics, the magnitude of the nEDM and the rates of proton and $0\nu\beta\beta$ nuclear decays induced by new operator structures must be calculated.




Figure 3-35. Observation of a nEDM would be evidence that our universe does not respect time-reversal symmetry. (Image available at Nature.com undated.)

The thrust of the fundamental symmetries effort heading into the exascale era is to provide the theoretical support necessary to capitalize on the experimental programs aimed at identifying and exposing the sources of BSM physics. At low energies, CP violation originating from BSM physics manifests itself through local CP-odd operators with coefficients that depend on the details of the BSM framework. The impact of these operators on the structure and interactions of the nucleon and light-nuclei will be calculated. Studies of candidate BSM scenarios have established the minimal precision goals of calculations required for nEDM calculations. In most scenarios, B violation induces the proton to decay, and DUNE will be relatively more sensitive to the $p \rightarrow \nu K$ decay mode. Observation of the $0\nu\beta\beta$ decay of nuclei would provide compelling proof of the violation of lepton number. To support the DOE's planned ton-scale $0\nu\beta\beta$ experimental program to search for such decays, theoretical calculations of the matrix elements of $0\nu\beta\beta$ -decay-inducing operators are required to be performed at the 10–20% level in light nuclear systems. This will enable constraints to be placed on hadronic-level operators that will then be used in nuclear many-body calculations.

If $0\nu\beta\beta$ decay or a nEDM is observed, or both, signaling the explicit violation of L or T, respectively, a future generation of detectors will be focused on more precisely constraining the structure of the underlying interactions to better isolate the responsible BSM physics. As the lattice QCD calculations in few-nucleon systems, that are anticipated to exist at that time, will be of modest precision, it will become necessary to launch a second generation of calculations in larger nuclear systems and with a more extensive operator structure to facilitate the desired refinement and identification of the origin of L or T violation.

Path to Exascale in Fundamental Symmetries and New Physics

- 2016: Calculate the dominant contributions to the nEDM.
- 2020: Estimate short-distance two-nucleon distance interactions contributing to $0\nu\beta\beta$ and nuclear EDMs and complete precision calculations of leading contributions to proton decay and to the rate of neutron–antineutron oscillations.
- 2025: Complete precision calculations of the two- and three-nucleon short-distance interactions contributing to $0\nu\beta\beta$ decay of nuclei and of interactions contributing to nEDM and nuclear EDMs.

3.4.2.5 Simulating the Quantum Vacuum: Generation of Gluon Field Configurations

A critical common element in addressing the science objectives in cold QCD with lattice QCD is the simulation of the strong interaction quantum vacuum (the results of a recent calculation are shown in Figure 3-36). More precisely, lattice QCD calculations of quantities of importance to cold nuclear physics involve sampling the zero-temperature correlations between the gluon fields at different points in space-time. In order to optimize scientific output, significant coordination among scientists is required in designing, producing, and analyzing these ensembles of gluon-field configurations that are generated through Markov Chain Monte Carlo techniques. We expect that this coordination will continue through the exascale era.




Figure 3-36. The quantum fluctuations of the gluon fields captured in a field configuration in one ensemble. The solid regions show enhanced action density and the vectors show orientation of a component of the gluon fields in space-time. (Image reproduced with permission from Derek Leinweber.)

The key algorithm for gauge generation, HMC algorithm and its variants, is a hybrid molecular dynamics (MD) Markov Chain method utilizing Hamiltonian MD in a fictitious MD-time with Metropolis Accept/Reject steps. Area-preserving and time-reversible MD integrators are used in order to satisfy detailed balance, which, along with ergodicity, is sufficient to correctly sample the equilibrium gluon-field probability distributions. At each step of the MD, solution of the lattice Dirac governing quark–gluon interactions is required, which is the dominant cost of current gauge generation. The solvers have traditionally been Krylov subspace iterative methods for sparse systems, such as conjugate gradients (CG) or stabilized biconjugate gradients (BiCGStab). The cost of these methods is driven by the condition number of the lattice Dirac operator, scaling inversely

with the quark mass, and has faced strong scaling issues in recent fat-node systems, such as GPU-based architectures. Recent developments have attempted to address both these issues: solvers with domain decomposed (DD) preconditioners have demonstrated superior scaling performance, and the inclusion of deflation and adaptive multi-grid (AMG) preconditioning promises to make the cost of simulations essentially independent of the quark masses.

Precision calculations in nuclear physics require the ensembles of gluon-field configurations to satisfy certain properties: (1) the quark masses in the calculation must match their physical values; (2) the physical lattice volumes need to be sufficiently large to avoid unwanted size and thermal effects; (3) multiple lattice spacings are required for controlled continuum extrapolations of observables; and (4) the effects of isospin breaking and electromagnetic interactions need to be included for certain observables. Because of the magnitude of present-day computational resources, current calculations cannot simultaneously satisfy these requirements and are typically carried out with unphysically heavy and degenerate light quarks and without the inclusion of electromagnetic effects. Only recently have isospin breaking and electromagnetism been introduced into U.S. simulation algorithms, but because of the resource constraints, they have been limited to relatively small volumes of space-time. An additional challenge expected in the next 10 years is a growth in the cost of generating independent configurations as the lattice spacing is reduced, making ensembles with finer resolution differentially more computationally expensive to generate than current simulations. The data needs of the generation are expected to be relatively modest compared to the required computational resources.


Path to Exascale in the Generation of Gluon Field Configurations

- **2016: Generate gluon configurations at the physical point with equal mass up and down quarks and without electromagnetism for use by all lattice QCD projects.**
- **2020: Generate multiple ensembles of gluon configurations at and near the physical point with large volumes and finer space-time discretizations.**
- **2025: Generate configurations including electromagnetism with a range of large volumes and discretizations. Complete algorithmic and software development required for performance-portable gauge generation.**

3.4.3 Cross-Cutting Research Directions

Addressing the U.S. cold QCD program requires dedicated efforts across a diverse range of disciplines. Currently, scientists are preparing application codes for readiness on the next generation of supercomputers, through the NERSC Exascale Application Program (NESAP) and Theta Early Science Program (Theta ESP) at ALCF, and through existing partnership with hardware makers: Intel and NVIDIA. In considering future systems, many of the challenges are cross-cutting and include discovering and implementing algorithms that can take advantage of the increases in parallelism offered by the new HPC architectures. Challenges are also faced in constructing portable software implementations that will be able to maintain good efficiency on diverse architectures. A common element in all the related research thrusts is computing the response of quarks moving in the gluon-field configurations. This movement is represented as the solution of large sparse linear system of equations, the Dirac equation, within a 4-D grid of space-time. The ambitious physics campaigns that are anticipated in the exascale era will require hundreds of thousands, or possibly millions, of such Dirac equation solutions combined into millions of separate terms for each ensemble. These calculations are expected to consume an increasingly

large portion of the computational resources dedicated to this program. A schematic of the present USQCD software stack that has been put in place to address these present-day challenges is shown in Figure 3-37.



The complexity of tensor contractions for elaborate matrix elements and nuclei is an increasing burden on the presently limited personnel available for software development. While this stage of present-day calculations is often accomplished in independent, high-throughput streams, the complexity of the resulting work flow involves challenging data management tasks that must often overcome serious I/O bottlenecks. The development of new algorithmic techniques, for instance, in multi-linear algebra, conducted in collaboration with computational and computer scientists, is essential to reduce the complexity of the contractions, and new software infrastructure is required to mitigate I/O demands. Partnerships between ASCR and NP in the area of scalable sparse linear solvers — including those that have components that can be tuned to emerging architectures, programming models, and performance portability and performance engineering— are essential. Specific challenges include overcoming on-node scaling limitations and writing code that is suitable for auto-tuning. Further, calculations with smaller lattice spacings and with larger volumes will require a revitalization of efforts in I/O, data management, and analysis strategies, which will likely benefit from colleagues with expertise in this area. Collaborations focusing on software engineering and software sustainability would also be valuable to ensure continuing high code quality going into the exascale era.

The DOE SciDAC program (DOE 2015), the recently enacted ECP (ECP.org undated), and the application readiness programs coordinated by the ASCR facilities are excellent vehicles for the necessary research and software development. These various collaborations also frequently fund regular staff and postdoctoral associates, which help both NP and ASCR with staff retention and workforce development. In cold QCD, a large fraction of its SciDAC3 funding has been used to effectively support laboratory-based researchers as the laboratories form centers of software development. Current partnerships between ASCR researchers and cold QCD domain scientists are shown in Figure 3-38. The recently funded lattice QCD ECP provides an important, and much-needed, expansion of the software effort in cold QCD. With this ECP, the laboratory focus is sharpened. There is scope for a similar project focused on university-based researchers for collaborations with applied mathematics and computer science groups. Critical to the success of the cold QCD program is the recruitment of talented young scientists in the areas of nuclear physics

and HPC. Scientists with a multidisciplinary background who can work comfortably in both areas are currently uncommon and, when identified, many find it difficult to pursue a viable career path. It is desirable to provide a reward structure for such scientists to allow both the recruitment and retention of high-caliber researchers at national laboratories and universities.

The analysis of the correlation functions to extract physically meaningful quantities is intrinsically statistical in nature through the sampling of the quantum vacuum. The time-dependent signal-to-noise ratio, and the higher moments of the distribution of correlation functions, presents analysis challenges at the forefront of areas of statistical research. The need for a complete quantification of all uncertainties associated with these calculations, and the size of associated computational resources that are required to obtain a specified level of precision, is making collaborations with statisticians increasingly important to successfully accomplishing the scientific goals in cold QCD.




Figure 3-38. The present-day linkages between ASCR researchers and domain scientists in cold QCD.
(Image reproduced with permission from Robert Edwards.)

There are major commonalities between the algorithmic research and software developments in lattice QCD in nuclear physics and lattice QCD in HEP. The HEP lattice QCD effort started with Wilson's work on formulating a lattice version of QCD, with major thrusts in spectrum of QCD and weak processes involving mesons. Much of the software used for cold lattice QCD calculations is currently developed in partnerships with HEP scientists and the algorithmic needs have significant overlap, leading to joint development teams. This close partnership between HEP and nuclear physics has proven critical to the evolution of the field, bringing lattice QCD into nuclear physics, and is expected to continue into the exascale era. Further, the capacity computing hardware deployed and operated by USQCD, used to address the distinct scientific objectives of both HEP and nuclear physics, is a joint effort between scientists in both communities.

3.4.4 Computing Needs and Requirements

Lattice QCD calculations in cold QCD for nuclear physics proceed through three distinct stages, as shown in Figure 3-39, each of which requires large-scale HPC resources, and a final data analysis phase, which is less computationally intensive but can utilize advanced data management technologies. The first phase of these calculations involves the generation of representative samples

of the the strong vacuum, known as gluon-field (or gauge-field) configurations. The second phase is the computation of the propagation of quarks in these fields to produce data objects known as “quark propagators.” The third phase, known as contractions, combines the quark propagators into correlation functions and physical observables. Each phase places a different emphasis on parallelism.




Figure 3-39. A schematic of the three-phase work flow currently used in cold QCD calculations for nuclear physics observables. (Image reproduced with permission from Robert Edwards).


Gauge configuration generation is a strong-scaling problem admitting only data parallelism and is responsible for roughly 25% of the computational cost of current nuclear physics calculations. Because of its strong-scaling needs and linear nature, it is ideally suited to run on long campaigns on capability supercomputers. In contrast, quark-propagator calculations can admit a large additional degree of ensemble-level parallelism. Propagators can be computed independently on each gluon configuration, and many propagators can be computed in parallel on each configuration. Hence, the most effective way to produce quark propagators is in high-throughput ensemble calculations. The majority of this phase of computation is spent in sparse linear solvers, which are bound in terms of on-node performance by memory bandwidth in the throughput regime, and also by communications interconnect bandwidth when running on many nodes in parallel. In the extreme strong-scaling regime with small problem sizes on each node, the bottlenecks also include latencies of the cache and memory subsystems and that of the interconnect fabric. The contraction stage reduces the 4-D data of the quark propagators, where the dimensions are the three space and the time dimension, to 1-D correlation functions in which the remaining dimension is the time dimension of the original lattice. Additional trivial parallelism can be exploited in the reduction of the 3-D spatial data because it can be performed independently for each value of the time dimension.

Accomplishing the cold QCD science objectives in the exascale era will require substantial computational investments. These estimates are based on (1) the aggregated needs of both the gluon field generation program needed to supply the initial data and of analysis projects identified

in Section 3.4.2 and (2) the cold QCD case studies. It is estimated that this program will need approximately 4.5 billion node hours in 2020, as measured in GPU node hours of OLCF's current Titan system (OLCF undated) and 17.2 billion node hours in 2025. The breakdown of these requirements between capability and capacity resources is shown in Figure 3-40. A modest need for growth of the memory per node in future systems is anticipated. Overall memory needs per run will grow from 8 TB today to more than 400 TB over 10 years. One primary factor in this growth is the increase in the number of sites in the space-time volumes required for the finer discretization of space-time, which drives other factors such as an increase in the size of tensor objects to be combined in the final contraction phase.

Meeting the target of spending less than 20% of the computational time performing I/O will require achieving a maximum I/O rate of around 100 GB/sec. It is anticipated that scratch-file system size needs (hot data) will grow to more than 8 PB over 10 years. This growth is primarily driven by the increasing throughput of jobs moving from the systems. Finally the permanent long-term storage needs and archival needs (cold data) will grow to around 31 PB and 315 PB, respectively, by 2025. The breakdown of these requirements between "hot data" and "cold data" is shown in Figure 3-40.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025 COLD QCD



* Exaflop-system-year refers to the total amount of computation produced by an exascale computer in 1 year.

Figure 3-40. The capability and capacity computing resource requirements (left panel) and the hot and cold data requirements (right panel) in 2025 to accomplish the science objectives of the Cold QCD program.

With the increasing data needs, I/O and data management will become increasingly challenging. Gluon-configuration generation produces relatively small amounts of data, and typically only about 5% of the calculation is spent performing I/O, while propagator and contraction codes are expected to spend approximately 20% of the runtime performing I/O. Further, the nature of the data changes between the three phases of calculation. Gluon-configuration generation provides large ensembles of lattices that are typically long lived and are akin to raw experimental data. They can have "hot" lifetimes during phases of propagator and correlation function generation and can remain archived in between such periods. Propagator calculations produce data that are too large to keep long term

and are typically discarded after the correlation functions have been produced in the contraction phase. Finally, the contraction phase produces data that are logically stored in databases for random access in the final stage data analysis. Given these factors, it is necessary to improve both that I/O and data management strategies, for example, by taking advantage of burst buffers. The final-stage data analysis needs focus on the selection of portions of correlation function data and fitting these to known models. This may be different from, and perhaps simpler than, some typical present-day data science scenarios of searching for patterns, or classifying data according to identified features. The expertise of ASCR data scientists with the algorithms and technologies in managing, organizing, and analyzing large-scale data can provide fresh insight into our methodology, and is a natural point for an exploratory collaboration to improve data management and analysis techniques currently used in cold QCD calculations.

Effective exploitation of architectural components in modern computer systems (caches, fast and slow memories, and interconnects) in a portable fashion across diverse architecture is challenging, and programs offered by ASCR facilities for application readiness such as NESAP, Theta ESP, and the Center for Accelerated Application Readiness (CAAR) are invaluable, especially because they each feature a performance portability component allowing multiple facilities to be targeted. Partnerships and relationships formed with facilities and their staff, such as consultants/catalysts in these programs, carry forward into production and into long-term algorithmic and performance development partnerships in a natural and organic way. Nuclear physics can also contribute to feeding back experiences to the facilities and other users through participation in workshops, code “hackathons” and training webinars.

Laboratory and university researchers play a crucial role in accomplishing the cold QCD science objectives. At present, those associated with the cold QCD effort are supported by institutional DOE and NSF grants and by SciDAC. Without these researchers, the scientific goals that have already been accomplished would not have been and the objectives in the exascale era will not be realized. We estimate that more than a doubling in the currently SciDAC-supported (or equivalent) researchers in the area of cold QCD is required to accomplish the objectives outlined in this and a previous (DOE SC 2016) report. It is anticipated that the recently supported ECP in lattice QCD will meet some of this need.

The capacity-computing hardware that is operated by USQCD through DOE support continues to be essential to the cold QCD program. In addition to providing half of the computing resources available to cold QCD in the United States, it has proven itself to be essential in “standing up” small projects and new ideas for algorithms and physics thrusts, and bringing them to the point at which they can be integrated into the large-scale physics program supported on leadership-class compute platforms. This is intimately related to the concept of the “speed of science,” in which having a component of the ecosystem dedicated to taking an idea from the moment of conception through to the first attempt at execution in a short time, is important for the project and in nurturing the curiosity and agility of junior scientists. The USQCD project has been at the forefront of supercomputer hardware design with the QCDSF, QCDOC that led to the Blue-Gene series of IBM machines, and it deployed the first large-scale GPU cluster (located at JLab). Needs have been identified for an expansion of the capacity-computing hardware available to the cold QCD program. The *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) contains an initiative to address this and other points.

To meet the challenges and realize the full scientific potential of current and future experiments, we require new investments in theoretical and computational nuclear physics. We recommend new investments in computational nuclear theory that exploit U.S. leadership in high-performance computing.

The cold QCD program has proven itself to be an outstanding arena in the training of junior scientists. This is because of the diverse skill sets that junior scientists have access to and can choose from, ranging from the domain sciences, to HPC, to algorithm creation, to software production, to work-flow design and, of course, the actual running on the largest supercomputers in the world. Junior scientists who have been trained within USQCD are successfully in career paths in universities, national laboratories and private industry, for example, data-science and related companies and hardware producers, such as IBM and NVIDIA. It is critical to increase the resources available for the training of junior scientists in this multidisciplinary environment that arises naturally on the road to accomplishing the scientific objectives of cold QCD. The training provided in this area is well matched to significant workforce needs of the national laboratories.

With the volume of data expected to be produced and retained, there is a growing need for implementing community-wide data-curation and data-provenance practices. The present team-based structure of the community, each working toward a number of science objectives, while sharing some of the intermediate results, such as gluon configurations and quark propagators, has led to the community-wide practice of noncompete agreements between scientists as a way to protect their intellectual property, consistent with data management plans. The data and results, in many instances, reside in repositories accessible by the larger community. It is desirable to have a system and set of protocols that allow usage agreements and data provenance to be openly and readily available to the community. Ideally, a system would evolve toward a complete tracking of data from creation, through verification and storage, through analysis, to publication. An integrated data genealogy that is accessible to the wider community would be of benefit.

Path to Exascale for Computational Cold QCD

- **2016:** In collaboration with ASCR/AM/CS, develop sparse system solvers and integrators to support generation of gluon configurations at the physical point for use by all lattice QCD projects. Develop scalable contraction work-flow methods for many-body systems. Search for exotic hadrons with many-body decays, providing guidance for GlueX.
- **2020:** Provide new work-flow support for detailed mapping of the 3-D structure of the nucleon, constraints on chiral nuclear forces for input into nuclear many-body calculations, and estimates of short-distance two-nucleon distance interactions contributing to $0\nu\beta\beta$ and nuclear EDMs.
- **2025:** In collaboration with ASCR/AM/CS, implement variance reduction techniques for precision matrix element computations of gluonic observables in hadronic and nuclear structure. These calculations will inform and guide the experimental program for the EIC and in fundamental symmetries.

3.5 Hot Quantum Chromodynamics

3.5.1 Scientific Challenges and Opportunities

If matter is heated to temperatures exceeding approximately 150 MeV (1.5×10^{12} K), it undergoes a transition from a system composed of hadrons (e.g., protons, neutrons, and pions) to a new state composed of liberated quarks and gluons; this new state is called the QGP. Such temperatures were generated in the early universe, and an analogous state of quarks and gluons, expected to exist at very high densities, may form the core of massive NSs and supernovae. To better understand these systems, it is important to (1) determine how matter subject to such extreme conditions behaves using first-principles QCD calculations, and (2) recreate it in a controlled, laboratory setting to study its properties. There has been tremendous experimental progress in the last 20 years using high-energy particle colliders, allowing scientists to recreate such extreme conditions at will. These capabilities have resulted in the generation, by the RHIC at Brookhaven and the LHC at CERN, of a wealth of experimental data concerning the behavior of QCD matter at high temperature and small net baryon density.

One of the most striking experimental findings concerns the collective flow of the matter generated in ultra-relativistic heavy-ion collisions. The magnitude of the observed flow and its dependence on transverse momentum, centrality, particle species, and other quantities, are consistent with viscous hydrodynamic models that describe the QGP as a nearly perfect fluid with a very low ratio of shear viscosity to entropy density ($4\pi\eta/s$ that lies between 1–3). Such low values of η/s were not expected based on leading-order perturbative QCD (pQCD) calculations, and this result has been taken as an indication of the strongly interacting nature of the QGP generated in ultra-relativistic heavy-ion collisions. Ongoing experiments are further probing our understanding of the energy scales at which a strong coupling (versus a weak-coupling) description of the QGP is necessary. The investments in experimental heavy-ion programs add up to several billion dollars. To capitalize on these investments, future developments in theory — in particular in computational theory — are essential.

Although our understanding of the physics of the QGP has increased significantly in the last decade, there are still many open questions that need to be addressed:

- What is the QCD phase diagram and EoS at finite density?
- How do the transport coefficients of strongly interacting matter depend on temperature and density?
- How are the early-time non-equilibrium dynamics of the QGP and its subsequent approach to thermal equilibrium realistically modeled?
- How can the deconfined nature of the new state of matter be experimentally established?
- Above what energy scale can the QGP be best described by weakly interacting quark and gluon quasiparticles?
- Are there hadronic excitations in the deconfined phase, and if so, what are their properties?

These open questions overlap with the key scientific questions that drive the field of nuclear physics, as identified in the *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015). As discussed in this Long Range Plan, the RHIC Beam Energy Scan program will take place in 2019–2020. The goal will be to study QCD matter as a function of density and search for the critical end point in the QCD phase diagram. This program will require an upgrade of the STAR detector (iTPC upgrade) that is already in progress. Clearly, lattice QCD studies of the phase diagram and EoS will be crucial for the interpretation of the experimental results. Beyond 2020, the goal of the RHIC program is to study properties of the QGP at various length scales through measurements of

penetrating probes, such as heavy quarkonia, jets, and open heavy-flavor mesons. The success of this experimental program will depend on performing first-principles calculations of the in-medium properties of mesons. A complementary program at the LHC will examine penetrating probes that will also be impacted by these same calculations.

The past 15 years of the RHIC experimental program and more than 5 years of the LHC heavy ion program have resulted in a wealth of experimental data on bulk particle production, and significantly more data are anticipated in the years to come. To optimally use these data to quantify QGP properties requires extensive numerical modeling by means of relativistic viscous hydrodynamics. In this context, one of the outstanding theoretical issues that is impeding our quantitative understanding of the properties of the hot and dense matter created in heavy-ion collisions is the description of the early-time dynamics in the collision systems prior to the achievement of approximate local thermal equilibrium. The early-time dynamics of these systems can be described using classical statistical calculations of the gluon fields. These calculations will provide the initial conditions needed for viscous hydrodynamic simulations.

3.5.2 Priority Research Directions

To meet the challenges of establishing a quantitative theoretical description of heavy-ion collisions and the properties of hot and dense QCD matter, progress in the following areas is required:

- Lattice QCD calculations at non-zero net baryon density;
- Lattice QCD calculation of the in-medium meson spectral functions and transport coefficients;
- Quantitative extraction of QGP properties through model-to-data comparison; and
- Classical statistical simulations of early time dynamics.

These priority research directions and their associated computational resource requirements are discussed in the following subsections.

3.5.2.1 Lattice QCD Calculations at Non-Zero Net Baryon Density

The use of modern supercomputers enables the calculation of fundamental properties of the QGP using first-principles lattice QCD calculations. Such calculations indicate that, at zero baryochemical potential ($\mu_B=0$), there is a smooth crossover between hadronic matter and the QGP at $T \sim 154$ MeV, with all thermodynamic quantities changing smoothly from one phase to the other (Bazavov et al. 2012[a]; Bhattacharya et al. 2014), as shown in Figure 3-41. This information is now being used as direct input into phenomenological models of particle production in heavy ion collisions because these require knowledge of the QCD EoS. Despite the success of the finite-temperature lattice QCD program, challenges related to the computation of the QGP's EoS at finite baryochemical potential remain. These challenges are related to the fact that at finite μ_B , the action becomes complex-valued and, as a result, standard Monte-Carlo-based algorithms to compute the path integral can no longer be used. To circumvent this problem, scientists can employ a Taylor expansion to expand the path-integral around $\mu_B = 0$ to a particular order in μ_B . Such an expansion can be used in the vicinity of $\mu_B \sim 0$ provided that the series has a non-vanishing radius of convergence; at present, most finite μ_B calculations compute expansion coefficients (susceptibilities) through the fourth order. Looking to the future, high-statistics computations using sixth-order expansions will require large-scale computational resources. Current calculations are performed on configurations with a relatively course lattice spacing. To obtain reliable results from an extrapolation to the space-time continuum, calculations at smaller lattice spacing are required, which also require significant computational resources.

An exascale ecosystem is necessary to further constrain the nature of the QGP phase transition at finite baryochemical potential. Model-dependent calculations suggest that there is a line of first-order phase transition in the μ_B - T plane, terminated by a critical end point. The search for this possible critical end point is one of the main goals of the ongoing RHIC BES experiment. In order to make reliable comparisons between the BES results and theoretical predictions, it is necessary to perform detailed hydrodynamic simulations of the QGP at finite baryochemical potential. For such comparisons, information about the EoS at finite μ_B is required. For this reason, a key focus of the finite temperature/density lattice QCD effort is the determination of the EoS at finite μ_B using high-order Taylor expansions of the partition function.




Figure 3-41. Finite-temperature and baryochemical potential lattice QCD calculations allow the exploration of the transition and crossover lines separating the low-temperature hadronic phase from the QGP phase. A summary diagram taken from the *2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015) is shown in the left panel. The left yellow-shaded band indicates the current constraints from lattice QCD, and the red points indicate the experimentally measured freeze-out curve (Bazavov et al. 2012[b]). The right panel shows the energy density minus three times the pressure, the pressure and entropy density of hot strongly interacting matter divided by the temperature to the power four calculated with two different lattice discretizations (stout and highly improved staggered quark [HISQ]) and extrapolated to the continuum (Bazavov et al. 2014).

Path to Exascale in Lattice QCD at Non-Zero Net Baryon Density

- **2016:** Perform calculations of thermodynamic quantities on lattices with temporal extent $N_t=8$.
- **2020:** Perform calculations of thermodynamic quantities on lattices with temporal extent $N_t=12$.
- **2025:** Perform calculations of thermodynamic quantities on lattices with temporal extent $N_t=16$ and obtain continuum extrapolated results.

3.5.2.2 Lattice QCD Calculations of In-Medium Spectral Functions and Transport Coefficients

The extraction of spectral functions from lattice QCD calculations will require significant computing resources. The study of penetrating probes, such as thermal photons and dileptons (Rapp 2011) and quarkonium (Mocsy et al. 2013), is critical to quantitatively understand the properties of matter produced in heavy ion collisions. In particular, the suppression of quarkonium production in heavy ion collisions, as shown in Figure 3-42, is considered an indication of QGP formation (Matsui and H. Satz 1986). One of the themes of the *2015 Long Range Plan for Nuclear Science* is probing the QGP at different length scales by studying quarkonium production in heavy-ion collisions. To interpret the experimental findings, the properties of quarkonia states at non-zero temperature (e.g., their thermal widths) are required, along with the disassociation temperature of different quarkonium states. These are encoded in the meson spectral functions, which can

be extracted from Euclidean time correlation functions calculable using lattice QCD. In general, hadron correlation functions are important for heavy ion phenomenology. Vector meson correlation and spectral functions in the light quark sector provide information on the thermal photon and dilepton rates, as well as on the electric conductivity of QGP. Many of the calculations of the meson spectral functions performed thus far have been in the quenched approximation to QCD; including dynamical quarks in future calculations will be necessary, and will necessitate an increase in computational resource requirements of two to three orders of magnitude.




Figure 3-42. Nuclear suppression factor R_{AA} for the $Y(1s)$ and $Y(2s)$ bottomonia states as a function of the number of nucleons participating in the heavy ion collision (N_{part}) (left) and particle transverse momentum (p_T) (right). Data are provided by the CMS and ALICE collaborations (CMS Collaboration 2015; Abelev et al. 2014). The curves are the predictions of a theoretical model that uses viscous anisotropic hydrodynamics to model the QGP evolution and takes into account the in-medium QGP decay width of the various Y states (Krouppa et al. 2015). The level of Y -suppression seen is much larger than that which would be generated by cold nuclear matter effects. This observation can be seen as an indication of QGP-induced suppression of bottomonia.

Path to Exascale in Lattice QCD Calculations of In-Medium Spectral Functions and Transport Coefficients

- 2016: Complete calculations of quarkonium spectral functions and transport coefficients in quenched approximation in the continuum limit.
- 2020: Achieve semi-realistic calculations of meson spectral functions and transport, including the effects of dynamical quarks.
- 2025: Complete fully realistic calculations of meson spectral functions and transport coefficients in the continuum limit.

3.5.2.3 Quantitative Extraction of QGP Properties through Model-to-Data Comparison

On the relativistic viscous hydrodynamics front, there are two areas that will require large-scale computing resources in the next decade: (1) event-by-event (EBE) 3+1-D viscous hydrodynamic simulations of QGP evolution (EBE vHydro) and (2) simulations of early-time, non-equilibrium dynamics using classical statistical gauge theory coupled to self-consistent solution of the Dirac equation. On the first front, the state-of-the-art modeling QGP evolution begins with Monte Carlo sampling of the initial energy-momentum tensor that is necessary to begin the evolution of the vHydro. This step is necessary because, on an EBE basis, there are fluctuations in the initial energy-density distribution that can cause, for example, a non-vanishing elliptic flow v_2 even in central collisions. In addition, EBE fluctuations induce higher-order azimuthal flows, such as triangular flow and quadrangular flow, that are quantified by the elliptic flow coefficients v_2 , v_3 , v_4 , and

higher. As a result, extracting quantitatively reliable values for the QGP transport coefficients, such as the shear and bulk viscosities, requires performance of thousands of EBE vHydro simulations coupled to a late-time hadronic cascade; the latter must be averaged over events using a statistical model for hadron production.

Recently, state-of-the-art statistical treatments have been applied to a global comparison of RHIC and LHC data with models, effectively constraining a high-dimension parameter space. These analyses have provided dramatic proof that experimental measurements from heavy ion collisions can be used to constrain the QCD EoS and transport coefficients (Pratt et al. 2015). In addition to sampling thousands of initial condition and freeze-out configurations, such statistical treatments require running the full model at thousands of separate points in parameter space to extract the necessary information from experimental data. These investigations will become more computationally demanding as scientists extend their viscous hydrodynamic codes to include finite baryochemical potential and fluctuations in the rapidity (longitudinal) direction.

Turning to the early-time dynamics of the QGP, as mentioned previously, modern viscous hydrodynamic simulations require EBE initial conditions. Currently, there are various models to describe this, some of which are more closely tied to QCD than others. The most basic model is called the Monte Carlo Glauber model; with this model, scientists construct an energy density profile based on sampling of the collisional cross-sections for “bags” of nucleons, depositing energy at the point of nucleonic collision. This model, while reasonably phenomenologically successful, has no firm basis in QCD, and uncertainties associated with its predictions cannot be quantified. To address this shortcoming, other types of initial conditions that take into account sub-nucleonic fluctuations have been developed in recent years. In particular, the impact-parameter-dependent Glasma model (IP-Glasma) has proved to be quite successful in describing the integrated elliptic flow and the full probability distribution of a number of flow coefficients. Results from the IP-Glasma model are shown in Figure 3-43.




Figure 3-43. Visualization of an IP-Glasma-based EBE initial condition for the energy density distribution generated in a high-energy heavy-ion collision (left), along with the evolved distribution at $\tau = 6$ fm/c using two values of the shear viscosity to entropy density ratio: $\eta/s = 0$ (middle) and $\eta/s = 0.16$ (right). (Images reproduced with permission from B. Schenke.)

Existing EBE viscous hydrodynamic codes need to be enhanced to include finite baryochemical potentials and longitudinal EBE fluctuations. More advanced dissipative hydrodynamics frameworks, such as anisotropic viscous hydrodynamics (see, for example, Bazow et al. 2014), are required for a more faithful description of early-time dynamics and dynamics in dilute regions. Such frameworks are necessary because the early-time QGP has relatively large deviations from thermal equilibrium, which can cast doubt on the application of standard viscous hydrodynamics approaches in certain regions of the QGP. The resulting viscous and anisotropic viscous hydrodynamics codes will also be tested in applications to small collision systems such as p+A and p+p. This testing will help to answer the question of what is the smallest drop of QGP that can be created in a laboratory setting. Present hydrodynamic simulations of heavy-ion collisions ignore the influence of thermal dynamical fluctuations, treating the hydrodynamic evolution module as a deterministic map of the EBE fluctuating initial conditions onto the final state. However, for dissipative systems, thermal fluctuations consistent with the non-zero values of the shear and bulk

viscosities of the fluid have to be taken into account during the dynamical evolution. Each sampled initial state should be evolved many times with the hydrodynamic code, to adequately sample the stochastic effects of thermal fluctuations on the dynamical evolution. This increases the numerical cost of the dynamical modeling of heavy-ion collisions by several orders of magnitude, pushing it far beyond present capabilities.

Path to Exascale in Quantitative Extraction of QGP Properties through Model-to-Data Comparison

- **2016:** Complete first statistical extraction of EoS and shear viscosity of QGP from heavy ion data.
- **2020:** Achieve precise extraction of QGP properties for small baryon density.
- **2025:** Achieve precise extraction of QGP properties at high baryon density.

3.5.2.4 Classical Statistical Simulations of Early Time Dynamics and Equilibration

The energy-momentum tensor resulting from the IP-Glasma model is computed by solving the classical Yang-Mills equations numerically and extracting the required color-averaged components. The use of classical Yang-Mills simulations is justified by the fact that the system produced in the wake of a high-energy heavy-ion collision is dominated by gluons because of the rapid growth in the gluon number at small- x . This rapid growth is related to the phenomenon of gluon saturation presented in the Color-Glass-Condensate (CGC) model (McLerran and Venugopalan 1994) which, in the wake of a heavy-ion collision, generates conditions lying between those of CGC and QGP, which is known as the Glasma (Gelis et al. 2009). Because of the high gluon occupation numbers in the initial state, the dynamics of the system can be calculated over some period of time, to good approximation, by using classical statistical gauge theory. In practice, this amounts to solving the real-time classical Yang-Mills equations in temporal axial gauge, making use of standard Hamiltonian-based algorithms to evolve the field configurations. While this program has helped to create more realistic models of the early-time dynamics of the QGP, to date researchers have not been able to faithfully model the production of quarks and anti-quarks in the Glasma background. This issue is complicated by the fact that classical simulations of fermion dynamics cannot be performed because their occupation numbers are bounded by the Pauli exclusion principle. As a result, researchers must solve the real-time Dirac equation in the presence of Glasma background fields; doing this reliably will require exascale computational facilities. The highly non-linear Yang-Mills evolution, which describes the non-equilibrium evolution of the saturated gluon field generated in high-energy collisions, is required and must be coupled to the Dirac equation in order to describe the production of quarks at early time after the nuclear impact. Such a description will provide essential inputs for viscous hydrodynamic calculations and insight into the chemical equilibration of the QGP.

Path to Exascale in Classical Statistical Simulations of Early Time Dynamics and Equilibration

- **2016:** Complete simulations in pure classical Yang-Mills theory, including exploratory calculations with dynamical fermions.
- **2020:** Complete simulations with dynamical fermions using Wilson formulations.
- **2025:** Complete fully realistic simulations with dynamical fermions using chiral fermions.

3.5.3 Cross-Cutting Research Directions

Achieving the goals of the U.S. hot QCD program requires efforts from a number of different disciplines, including analytic QCD, HEP, and computational physics. To date, development of the computational software that researchers employ has been accomplished primarily by the domain scientists. Domain scientists working on lattice QCD calculations at non-zero baryon chemical potential are involved in the NESAP and have non-disclosure agreements with vendors such as Intel and NVIDIA. Looking to the future, the need to fully optimize all hot QCD codes will require an increasing involvement from scientists supported by ASCR, particularly for memory optimization on GPU architectures.

Like zero-temperature lattice QCD applications, all three of the lattice-based components discussed here require the solution of a large system of sparse linear equations to obtain solutions to the Dirac equation. This is true for the determination of the QCD EoS, determination of the spectral functions in unquenched configurations, and the inclusion of quark-antiquark production at early times after a heavy-ion collision. Algorithmic improvements at this level would have important cross-cutting benefits. In addition, algorithmic improvements for lattice calculations will have benefits for both the hot and cold QCD programs, as well as for lattice QCD calculations in HEP because there is a great deal of synergy between these research directions. This synergy is reflected in the mission of the USQCD collaboration, whose near-term goals are as follows:

1. Calculate the effects of strong interactions on weak-interaction processes and search for observables that are sensitive to BSM physics.
2. Determine the properties of strongly interacting matter subject to extreme conditions such as those created in relativistic heavy-ion collisions.
3. Calculate the masses of strongly interacting particles and obtain a quantitative understanding of their internal structure.
4. Lay the foundations for investigations of strongly interacting sectors of new physics that may be discovered at the LHC.

We also note that the real-time lattice QCD codes, which compute early-time quark/anti-quark production in a self-consistent CGC background, could benefit greatly from increased interactions with both hot and cold lattice QCD scientists, who can help to optimize codes using the wealth of experience accumulated in this application domain.

Finally, the hydrodynamics applications discussed here are quite distinct in their requirements from other hot QCD applications, which traditionally emphasize lattice QCD. However, in the hydrodynamics domain, there is potentially a great deal of overlap with groups who are performing supernovae simulations, because both rely on solving 3+1-D hydrodynamics and kinetic theory on large space-time lattices. Looking forward, the hydrodynamics applications would benefit from having the codes ported to GPUs and many-core architectures, because the algorithms used to solve the necessary partial differential equations are local in space-time. It would be useful to have input from computer scientists who have experience with these architectures and, most importantly, with memory optimization.


In the short term, progress on all fronts is hindered by a lack of qualified personnel to work on the various projects outlined here. The funded “pipeline” for the recruitment and retention of faculty and postdoctoral associates needs to be enhanced. In the realm of long-term program support, it is important to have a career path for computational physicists. In order to foster ongoing collaboration with ASCR, we recommend that these programs consider hiring domain scientists at leadership-class computing facilities in the area of hot QCD who would contribute to our long-term goals.

3.5.4 Computing Needs and Requirements

Hot QCD research currently makes use of a large amount of capability and capacity computational resources. Most of the cycles are spent on performing lattice QCD calculations of thermodynamic quantities at non-zero chemical potentials, as well as meson spectral functions. These calculations presently use approximately 200 M Titan (equivalent) core hours and 150 M Blue Gene/Q core hours, respectively. The computational resources expended on modeling heavy-ion collisions and classical statistical simulations of early-time dynamics are presently an order of magnitude smaller — about 15 M core hours. In the next decade, these calculations will require significantly more computing resources.

Calculations of the thermodynamic quantities at non-zero chemical potentials have to be performed on finer lattices, increasing the computational needs by a factor of approximately 300. The calculations of meson spectral functions should include the effects of dynamical quarks, increasing the time requirements by three orders of magnitude. Performing classical statistical simulations with dynamical fermions on sufficiently large lattices will also require a significant increase in computer cycles. Finally, hydrodynamic modeling of heavy ion collisions should rely on 3-D hydrodynamics simulations for multiple collisions energies and fluctuating initial conditions, again increasing the requirements for computational time by two orders of magnitude. We have summarized both anticipated system compute hours and data storage needed for the Hot QCD program in Figure 3-44.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025 HOT QCD



* Exaflop-system-year refers to the total amount of computation produced by an exascale computer in 1 year.

Figure 3-44. The capability and capacity computing resource requirements (left panel) and the hot and cold data requirements (right panel) in 2025 to accomplish the science objectives of the Hot QCD program.

Path to Exascale for Computing Needs and Requirements

- 2016: In collaboration with ASCR and applied mathematics and computer science experts, begin porting codes for viscous hydrodynamics and classical statistical simulations to GPU platforms and other many-core architectures and commence working on new work flow models to calculate thermodynamic quantities at non-zero baryon density.
- 2020: Perform large-scale calculations aimed at quantifying the properties of QGP from data to experiment using optimized code on pre-exascale platforms. Perform realistic classical statistical simulations of early-time dynamics using Wilson fermions on pre-exascale platforms. Perform lattice QCD calculations at non-zero baryon density on $N_t=12$ lattices.
- 2025: Carry out the exascale calculations needed for successful completion of the RHIC mission: perform fully realistic lattice QCD calculations of thermodynamic properties at non-zero density and spectral functions in the continuum limit. Quantify QGP properties in a wide range of baryon density from model to experiment comparison. Study early time dynamics in heavy collisions using chiral fermions.

4 PATH FORWARD

This review addressed the extensive, diverse, and evolving requirements of the exascale HPCDE required to meet the scientific challenges of the DOE and NSF nuclear physics research programs. The diversity of the nuclear physics HPCDE emerges from the need to solve QCD in the hadronic and low-energy regimes and at extremes of temperature and density, calculate the structure and reactions of nuclear many-body systems, perform large-scale nuclear astrophysics simulations, and acquire and analyze large data sets generated in the nation's laboratories using real-time data-processing capabilities. This HPCDE includes the following:

- Exascale capability computing resources comprising heterogeneous architectures with commensurate hot- and cold-data storage capabilities;
- Enhanced infrastructure for data-handling capabilities — including transfer, archiving, curation, and large-scale databases;
- Development of a highly skilled and sustainable workforce, including the creation of permanent positions at the interface between nuclear physics and applied mathematics and computer science; and
- Increased capacity computing resources at universities and nuclear physics laboratories.

Increased investment in each of these components of the ecosystem is required to allow nuclear physicists to optimally address the scientific challenges using exascale-computing resources. The integration of an exascale HPCDE into the nuclear physics research program, in coordination with advances in nuclear theory, experiment, applied mathematics, statistics, and computer science, will provide unprecedented predictive capabilities in subatomic physics and is assured to spark scientific discoveries that we cannot imagine today.

4.1 Designing an Exascale Ecosystem for Nuclear Physics

For researchers to successfully achieve the scientific goals described in this report, they need access to a rich computing and data ecosystem. This joint NP-ASCR review identified items of importance to the NP community. Scientists performing research of interest to NP need to (in no particular order):

- Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales;
- Explore parameter spaces and quantify uncertainties using ensembles of calculations;
- Verify and validate algorithms, codes, and models;
- Develop and optimize new codes, algorithms, models, and workflows for exascale computing architectures;
- Read, write, manage, analyze, curate, and track data of complexity and scale never before encountered;
- Grow and sustain a workforce to carry NP computational science through the exascale era, including enhanced collaboration between NP, ASCR, and NSF, and creation of permanent positions at the interface between NP and ASCR; and
- Enhance local capacity hardware at laboratories and universities that is capable of efficiently integrating with ASCR infrastructure and facilities.

Each of these tasks is described in the paragraphs that follow.

Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales

The computational resources required during the next decade by nuclear physicists have been established. Figure 4-1 shows the requests for compute resources during 2016, along with projected resource requests for 2020 and 2025. All areas except Experiment and Data have very large requirements for heterogeneous resources, either GPU or many-core. Experiment and Data and all areas have modest projected capacity computing needs for conventional core resources. By 2025, nuclear physics is projected to need exascale data storage for both cold- and hot-data at the levels shown in Figure 4-1, and the associated infrastructure to transfer these data for subsequent use and analysis. Nuclear physics research has a clear need for both exascale capability computing resources and exascale capacity computing resources and the associated global data-handling infrastructure. The nature of some of the calculations required to be performed in the ecosystem have large in-node memory needs, along with the need for fine control over memory hierarchies. Further, these nodes will be required to have high-speed I/O capabilities.




Figure 4-1. The projected computational resources required by nuclear physicists. The left panel shows the needs for conventional CPU (millions of core hours per year) and heterogeneous architectures (billions of titan node hours per year). The right panel shows the data storage needs for hot data and cold data. The opacity of the symbol denotes the year (either 2016, 2020, or 2025). [Note: the Hot QCD conventional CPU expected resource requests differ from those stated in the white paper because it is assumed that future collaborations will permit code porting to heterogeneous architectures.]

Computing at exascale requires access to exascale hardware, supporting software, and applications that can run efficiently in an exascale environment. ASCR facilities already host pre-exascale systems and, in conjunction with domain scientists (NESAP at NERSC, CARR at OLCF, and ESP at ALCF), they have developed successful and effective application readiness programs and will continue to do so. ASCR facilities will host exascale systems when they become available, with configurations, options, and engineering-developed capabilities informed by the needs and requirements laid out in the Exascale Requirements Review reports.

The facility application readiness programs are focused on achieving performance of specific algorithms on specific hardware architectures. But that is not enough to take full advantage of the opportunities afforded by exascale. New algorithms and mathematical formulations will be needed, and new computing and runtime paradigms need to be explored. There is great need in all areas of nuclear physics for enhancing the collaboration between NP domain scientists and ASCR applied mathematicians, statisticians, computer scientists, and facility experts – along with HEP,

BES, and other physicists. The goal of such collaborations, which will build on and enhance the very successful SciDAC collaborations, will be to develop codes and work flows to operate in an exascale environment.

Explore parameter spaces and quantify uncertainties using ensembles of calculations

Modern HPCDEs give scientists the ability to explore the appropriate parameter space in calculations and simulations. This capability is required in order to, for example, test theories and assumptions, find optimal solutions, propagate uncertainties, validate models, and quantify uncertainties. These endeavors require access to large numbers of core-hours distributed across a spectrum of job scales on HPCDEs. Allocations of adequate size — often mid-scale computing allocations — are needed to support what are often referred to as “ensembles” that are run to enable solution of exascale-size challenges with a complete quantification of associated uncertainties.

Ensemble calculations and simulations, and the experimental program, will require rapid-access and addressable databases. Such databases will also be required for data curation and long-term effective staging of hot and cold data for post-production analysis. Currently, database are integrated into some production codes; an enhanced effort, through collaboration with ASCR in developing databases suitable for conducting nuclear physics in the exascale era, is now required.

Verify and validate algorithms, codes, and models

On-demand access to hardware of the same or similar architecture as the exascale resources is vital for code development, testing, and scaling studies. Enhanced coordination between NP, NSF, and ASCR facilities in planning for the requirements of nuclear physics would be beneficial in assuring that such resources are built into procurement and NSF, ERCAP, INCITE, and ALCC allocation planning. Capacity computing resources at universities and national laboratories have traditionally provided the infrastructure on which such development and validation occur, primarily because of the “speed of science” and the ability to gain access without the typically long-delays associated with peer review and committee evaluation.

Develop and optimize new codes, algorithms, models, and work flows for exascale computing and architectures

Creating and implementing efficient, high-performance scientific applications, and the work flows used to execute them at exascale, are critical to accomplishing the science objectives of nuclear physics. Of increasing importance, as computer architectures evolve, is the development of abstractions and tools for performance portability. Developing close collaborations between nuclear physics researchers and ASCR researchers and staff at the facilities and building upon those that have been set in motion with the SciDAC projects, is essential to achieving this level of computational sophistication.

Read, write, manage, analyze, curate, and track data of complexity and scale never before encountered

Issues related to data, and data science, are now at the forefront of HPC, and this trend is expected to continue in the exascale era. These issues include (1) integration with existing (non-HPC) work flows and methods for data acquisition, transfer, and analysis in order to leverage the scale of resources available at HPC centers, and (2) working seamlessly across computational and data environments at all scales and locations relevant to nuclear physics.

As calculations and simulations reach exascale, the time needed to produce output data and checkpoint files with current technology and algorithms grows to unacceptable levels. A close collaboration between ASCR and the nuclear physics community can address this issue through development of algorithms and technologies like burst buffers and identification of system balances that accommodate the required I/O rates.

Experimental nuclear physics science is embracing HPC to solve its significant needs for additional computational resources and extreme data acquisition, storage, sharing, management, and analysis. The nuclear physics experimental community has a mature and widely accepted software, data management and movement, and analysis infrastructure that will have to interoperate with HPC facilities. Significant effort will be required to prepare nuclear physics experimental software and work flows and develop portable software stacks for HPC. Such efforts will also be required to accommodate existing and evolving nuclear physics infrastructure at the HPC facilities and the fading distinction between online and offline analysis. Moving forward, ASCR, NSF, and NP scientists working together can help ensure that state-of-the-art analysis techniques are employed and are able to run efficiently at ASCR and NSF HPC facilities.

Further, there are emerging needs for enhanced real-time data capabilities. Deploying appropriate computational hardware at the nuclear physics experimental facilities that have architectures similar to those of ASCR facilities may be of benefit in interfacing nuclear physics experimental data requirements with ASCR HPCDEs.

First-of-their-kind HPC and data systems are — by their very nature — unique, with special environments. The diversity across ASCR and NSF HPC facilities — each with their special environments, requirements, and practices — can make it difficult for researchers to operate across facilities. But by the facilities working with each other and the science community, the key obstacles can be addressed and minimized. Examples include implementing common identification and authorization mechanisms, defining and implementing a standard base set of HPC software and tools, standardizing procedures and namespaces for using data transfer nodes, and developing a common set of tools to collate data with their corresponding data management plans.

There is a need for enhanced collaboration between ASCR and NP in data management across all areas of nuclear physics, including database optimization, coordination with approved data management plans, access control, data tracking from creation to deletion, and data curation.

Grow and sustain a workforce to carry NP through the exascale era, including enhanced collaboration between NP, ASCR, and NSF and creating permanent positions at the interface between NP and ASCR

The SciDAC projects provided the organization within the nuclear physics community to optimize the science output from tera- and petascale compute resources. The funding from NP and ASCR was used to support junior scientists — typically postdoctoral fellows and graduate students and fractions of senior scientists — in collaborating to develop new algorithms and optimize codes on evolving capability and capacity computing hardware. Without the SciDAC programs, the nuclear physics community would not have codes running efficiently on, for instance, GPU and many-core architectures nor on the Blue-Gene series. It is critical that such support for scientists and such collaborations continue through the exascale era. As identified in *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science* (DOE and NSF 2015), the Nuclear Physics Workforce Report (NSAC 2014[b]), and the National Academies nuclear physics decadal study report (National Academies Press 2013), a significantly larger HPCD-skilled workforce is required than exists today to meet the grand challenges facing nuclear physics.

Computational scientists in NP require significant support in developing and optimizing codes, algorithms, and work flows for next-generation systems. Building upon the developments and organization within the nuclear physics community that have emerged from the SciDAC projects, the NESAP program at NERSC, CARR at OLCF, ESP at ALCF, and ECP have initiated this process. Lessons learned and the best practices from these programs are being disseminated to DOE and NSF HPC users, in general, through the facilities training programs. All these efforts are valuable and are expected to continue. In addition, general training in HPC is being carried out,

with the Argonne Training School on Extreme-Scale computing being the current premier example. The ASCR facilities will team with the ECP to continue and expand programs like this.

Establishing viable career paths to facilitate the transition of junior scientists to senior scientists who remain associated with nuclear physics research efforts — either at universities or at the national laboratories — is becoming urgent. The larger question of workforce development, achieved through programs for career development, recognition, acceptance, and placement, is recognized as a major concern for science in general.

The major areas of collaboration that have proven fruitful for nuclear physics have been with high-energy, condensed matter, and atomic physicists and applied mathematicians and computer scientists. These collaborations must continue and become significantly stronger through the exascale era. In addition, collaborations with statisticians are forming, and they are expected to become essential considering the need for uncertainty quantification in theoretical computations and the dramatic increases in data volume and complexity that are expected in the exascale era.

Enhance local capacity computing and associated infrastructure at laboratories and universities that is capable of efficiently integrating with ASCR infrastructure and facilities

Data acquisition, processing, and analysis, which are key elements of the nuclear physics experimental program, require a diverse exascale ecosystem. Significant elements of the ecosystem — those required to acquire data from the detector, apply acceptance criteria, rapidly move accepted events to disk, and format the data for distribution to scientists — is accomplished at the national laboratories on local capacity-computing hardware. In the exascale era, petascale capacity-computing resources will be required at the national laboratories to accomplish the tasks that are necessarily local to the experimental detectors. Comparable-scale resources will be required at universities to perform the subsequent analyses and code optimization; such resources are critical to the education of junior scientists who will form the future workforce.

Capacity-computing resources at universities and national laboratories have proven invaluable in developing the codes and algorithms used for large-scale simulations and calculations in nuclear physics. In addition, such resources permit exploration of new ideas at a modest scale without the need for writing a proposal to be peer reviewed and possibly supported — a process that typically takes months. The nuclear physics long-range plan (DOE and NSF 2015) has identified a need for enhanced mid-scale capacity-computing hardware at universities and national laboratories to meet these increasing needs.

4.2 Summary

This review has identified essential elements of an exascale ecosystem that are required to accomplish the scientific grand challenges in nuclear physics. The ecosystem will have to be multi-faceted and multi-scaled. Diverse arrays of computational hardware are required, from exascale capability and capacity computing located at the ASCR and NSF facilities, to capacity computing located at the relevant national laboratories and universities. Significant growth of a diverse nuclear physics work force, working in close collaboration with ASCR and NSF scientists, is required to optimize the scientific accomplishments through the exascale era. There is a differentially increasing need for ASCR and NSF support in all areas related to nuclear physics data.

Working together with ASCR and NSF facilities and experts, nuclear physics research is going to be changed in remarkable ways.

5 REFERENCES

- Abelev, B.B., 2014, et al. (ALICE Collaboration), “Suppression of $\Upsilon(1S)$ at Forward Rapidity in Pb–Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” *Physics Lett. B* 738, 361.
- Accardi, A., J.L. Albacete, M. Anselmino, et al., 2012, “Electron Ion Collider — The Next QCD Frontier: Understanding the Glue that Binds Us All,” *Eur. Phys. J. A* 52 (2016) no.9, 268.
- Bazavov, A., T. Bhattacharya, M. Cheng, C. DeTar, H.-T. Ding, Steven Gottlieb, R. Gupta, P. Hegde, U. M. Heller, F. Karsch, E. Laermann, L. Levkova, S. Mukherjee, P. Petreczky, C. Schmidt, R. A. Soltz, W. Soeldner, R. Sugar, D. Toussaint, W. Unger, and P. Vranas, 2012[a], “Chiral and Deconfinement Aspects of the QCD Transition,” March 5, *Phys. Rev. D* 85, 054503.
- Bazavov A., H.-T. Ding, P. Hegde, O. Kaczmarek, F. Karsch, E. Laermann, S. Mukherjee, P. Petreczky, C. Schmidt, D. Smith, W. Soeldner, and M. Wagner, 2012[b], “Freeze-Out Conditions in Heavy Ion Collisions from QCD Thermodynamics,” November 9, *Phys. Rev. Lett.* 109, 192302.
- Bazavov, A., T. Bhattacharya, C. DeTar, H.-T. Ding, S. Gottlieb, R. Gupta, P. Hegde, U.M. Heller, F. Karsch, E. Laermann, L. Levkova, S. Mukherjee, P. Petreczky, C. Schmidt, C. Schroeder, R.A. Soltz, W. Soeldner, R. Sugar, M. Wagner, and P. Vranas, 2014, “Equation of State in (2+1)-flavor QCD,” November 4, *Phys. Rev. D* 90, 094503.
- Bazow, D., U. Heinz, and M. Strickland, 2014, “Second-Order (2+1)-Dimensional Anisotropic Hydrodynamics,” November 21, *Phys. Rev. C* 90, no. 5, 054910.
- Beane, S.R., E. Chang, S.D. Cohen, W. Detmold, H.W. Lin, T.C. Luu, K. Orginos, A. Parreño, M. J. Savage, and A. Walker-Loud (NPLQCD Collaboration), 2013, “Light Nuclei and Hypernuclei from Quantum Chromodynamics in the Limit of SU(3) Flavor Symmetry,” *Phys. Rev. D* 87(3):034506.
- Beane, S.R., E. Chang, S. Cohen, W. Detmold, H.W. Lin, K. Orginos, A. Parreño, M.J. Savage, and B. C. Tiburzi (NPLQCD Collaboration), 2014, “Magnetic Moments of Light Nuclei from Lattice Quantum Chromodynamics,” *Phys. Rev. Lett.* 113(25):252001.
- Beane, S.R., E. Chang, W. Detmold, K. Orginos, A. Parreño, M.J. Savage, and B.C. Tiburzi (NPLQCD Collaboration), 2015, “Ab Initio Calculation of the $np \rightarrow d\gamma$ Radiative Capture Process,” *Phys. Rev. Lett.* 115 (13):132001.
- Bertsch, G.F., D.J. Dean, and W. Nazarewicz, 2007, “Computing Atomic Nuclei,” *SciDAC Review*, 6, 42, <http://web.archive.org/web/20130530183059/http://www.scidacreview.org/0704/html/unedf.html>, accessed January 2017.
- Bhattacharya, T., M.I. Buchoff, N.H. Christ, H.-T. Ding, R. Gupta, C. Jung, F. Karsch, Z. Lin, R.D. Mawhinney, G. McGlynn, S. Mukherjee, D. Murphy, P. Petreczky, D. Renfrew, C. Schroeder, R.A. Soltz, P.M. Vranas, and H. Yin, 2014, “QCD Phase Transition with Chiral Quarks and Physical Quark Masses,” February 21, *Phys. Rev. Lett.* 113, no. 8, 082001.
- BNL (Brookhaven National Laboratory), undated, “Relativistic Heavy Ion Collider,” <http://web.archive.org/web/20170101132536/http://www.bnl.gov/RHIC/>, accessed September 2016.
- Bruenn, S.W., E.J. Lentz, W.R. Hix, A. Mezzacappa, J.A. Harris, O.E. Bronson Messer, E. Endeve, J.M. Blondin, M.A. Chertkow, E.J. Lingerfelt, P. Marronetti, and K.N. Yakunin, 2016, “The Development of Explosions in Axisymmetric Ab Initio Core-Collapse Supernova Simulations of 12–25 $M_{\odot}$ Stars,” *Astro. J. Lett.*, Volume 818, Number 2, 123.
- Calder, A.C., B.K. Krueger, A.P. Jackson, and D.M. Townsley, 2013, “The Influence of Chemical Composition on Models of Type Ia Supernovae, *Frontiers of Physics*, Volume 8, pp. 168–188. DOI: 10.1007/s11467-013-0301-4.

California Institute of Technology, undated, “Laser Interferometer Gravitational-Wave Observatory,” <http://web.archive.org/web/20161231065446/https://www.ligo.caltech.edu/>, accessed September 2016.

Caplan, M.E., A.S. Schneider, C.J. Horowitz, and D.K. Berry, 2015, “Pasta Nucleosynthesis: Molecular Dynamics Simulations of Nuclear Statistical Equilibrium,” *Phys. Rev. C* 91, No. 6, 065802, arXiv:1412.8502.

Carlidge, E., 2015, “Billion-Dollar Particle Collider Gets Thumbs Up,” *Nature* 521(7552):272, DOI 10.1038/521272a, <http://web.archive.org/web/20161002015621/http://www.nature.com/news/billion-dollar-particle-collider-gets-thumbs-up-1.17579>, accessed September 2016.

Cherry, J., G.M. Fuller, J. Carlson, H. Duan, and Y.-Z. Qian, 2010, “Multi-Angle Simulation of Flavor Evolution in the Neutronization Neutrino Burst from an O-Ne-Mg Core-Collapse Supernova,” *Phys. Rev. D* 82, 085025.

Cherry, J., J. Carlson, A. Friedland, G.M. Fuller, and A. Vlasenko, 2012, “Neutrino Scattering and Flavor Transformation in Supernovae,” *Phys. Rev. Lett.*, Volume 108, 261104.

Cirigliano, V., G.M. Fuller, and A. Vlasenko, 2015, “A New Spin on Neutrino Quantum Kinetics,” *Phys. Lett. B*, Volume 747, pp. 27–35.

CMS Collaboration, 2015, “CMS Heavy-Ion Public Physics Results,” CMS-PAS-HIN-10-006, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsHIN15001>, accessed September 2016.

DOE (U.S. Department of Energy), 2009, *Scientific Grand Challenges: Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale*, sponsored by NP and ASCR, January 26–29, https://web.archive.org/web/20150910200653/http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Np_report.pdf, accessed September 2016.

DOE, 2014, *Building for Discovery: Strategic Plan for U.S. Particle Physics in the Global Context, Report of the Particle Physics Project Prioritization Panel (P5)*, May, http://science.energy.gov/~media/hep/hepap/pdf/May-2014/FINAL_P5_Report_053014.pdf, accessed September 2016.

DOE, 2015, “SciDAC: Scientific Discovery through Advanced Computing,” last updated August 31, <http://web.archive.org/web/20160902031832/http://www.scidac.gov/>, accessed September 2016.

DOE and NSF (U.S. Department of Energy and the National Science Foundation), 2015, *Reaching for the Horizon: The 2015 Long Range Plan for Nuclear Science*, October, DOE Office of Science, Office of Nuclear Physics and NSF Division of Physics, Nuclear Physics Program http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf, accessed September 2016.

DOE Leadership Computing, 2016, “INCITE Program: Leading the Way for Scientific Computing,” <http://web.archive.org/web/20161224231904/http://www.doeleadershipcomputing.org/incite-program/>, accessed September 2016.

DOE SC (U.S. Department of Energy Office of Science), 2016, *ASCR Workshops and Conferences, Scientific Grand Challenges Workshop Series*, March 5, <http://science.energy.gov/ascr/community-resources/workshops-and-conferences/grand-challenges/>, accessed September 2016.

DOE SC, undated, “Nuclear Science Advisory Committee (NSAC),” <http://web.archive.org/web/20160809233306/http://science.energy.gov/np/nsac>, accessed September 2016.

- Dohet-Eraly, J., P. Navrátil, S. Quaglioni, W. Horiuchi, and F. Raimondi, 2016, “ $^3\text{He}(\alpha, \gamma)^7\text{Be}$ and $^3\text{H}(\alpha, \gamma)^7\text{Li}$ Astrophysical S Factors from the No-core Shell Model with Continuum,” *Phys. Lett. B* 757, 430.
- Duan, H., A. Friedland, G. McLaughlin, R. Surman, 2011, “The Influence of Collective Neutrino Oscillations on a Supernova r Process,” *J. Phys. G: Nucl. Phys.* 38, 035201.
- Duan, H., G.M. Fuller, and J. Carlson, 2008, “Simulating Nonlinear Neutrino Flavor Evolution,” *Computational Science and Discovery*, Volume 1, Number 1, 015007.
- Duan, H., G.M. Fuller, and Y.-Z. Qian, 2010, “Collective Neutrino Oscillations,” *Annual Review of Nuclear and Particle Science*, Volume 60, pp. 569–594. DOI: 10.1146/annurev.nucl.012809.104524.
- Dudek, J.J., R.G. Edwards, P. Guo, and C.E. Thomas (for Hadron Spectrum Collaboration), 2013, “Toward the Excited Isoscalar Meson Spectrum from Lattice QCD,” *Phys. Rev. D* 88(9):094505.
- ECP.org, undated, “Building a Capable Exascale Ecosystem,” DOE and NNSI (National Nuclear Security Administration), <https://web.archive.org/web/20161212070342/https://exascaleproject.org/>, accessed December 2016.
- Erler, J., N. Birge, M. Kortelainen, W. Nazarewicz, E. Olsen, A.M. Perhac, and Mario Stoitsov, 2012, “The Limits of the Nuclear Landscape,” *Nature*, 486; 7404.
- ESNet, undated, “The Network: A Nationwide Platform for Science Discovery,” <https://www.es.net/engineering-services/the-network/>, accessed October 2016.
- Fermilab, 2016, “Deep Underground Neutrino Experiment,” <http://www.dunescience.org>, accessed September 2016.
- Fermilab, undated[a], “Neutrinos,” <https://web.archive.org/web/20160827113410/http://www.fnal.gov/pub/science/particle-physics/experiments/neutrinos.html>, accessed September 2016.
- Fermilab, undated[b], “Facility for the Deep Underground Neutrino Experiment” <https://web.archive.org/web/20160729031224/http://lbnf.fnal.gov/>, accessed September 2016.
- FRIB (Facility for Rare Isotope Beams), 2012, *FRIB: Opening New Frontiers in Nuclear Science: Moving Forward with the Long Range Plan*, FRIB Users Organization for the NSAC Long Range Plan Implementation Subcommittee, August, https://web.archive.org/save/_embed/http://frib.msu.edu/_files/pdfs/frib_opening_new_frontiers_in_nuclear_science.pdf, accessed September 2016.
- FRIB, undated, “Facility for Rare Isotope Beams at Michigan State University,” <https://web.archive.org/web/20160903021107/http://frib.msu.edu/>, accessed September 2016.
- Gade A., and R. Zegers (eds.), 2014, *HRS: A High Rigidity Spectrometer for FRIB*, December, <https://web.archive.org/web/http://hrs.lbl.gov/documents/HRS-WhitePaper.pdf?attredirects=0&d=1>, accessed January 2017.
- Gandolfi, S., J. Carlson, and S. Reddy, 2012, “Maximum Mass and Radius of Neutron Stars, and the Nuclear Symmetry Energy,” *Phys. Rev. C* 85 (3), 032801.
- Gelis, F., T. Lappi, and L. McLerran, 2009, “Glittering Glazmas,” *Nucl. Phys. A*, Volume 828, 149, September.
- Helland, B., 2016, *NP Requirements Review*, presentation at the Nuclear Physics Exascale Requirements Review Meeting, U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, Gaithersburg, Maryland, June 15.

Hergert, H., S.K. Bogner, T.D. Morris, A. Schwenk, and K. Tsukiyama, 2016, “The In-Medium Similarity Renormalization Group: A Novel Ab initio Method for Nuclei,” *Phys. Rept.* 621, 165.

Jacobs, A.M., M. Zingale, A. Nonaka, A.S. Almgren, and J.B. Bell, 2016, “Low Mach Number Modeling of Convection in Helium Shells on Sub-Chandrasekhar White Dwarfs II: Bulk Properties of Simple Models,” June 8, arXiv:1507.06696v2 [astro-ph.HE].

JLab (Thomas Jefferson National Accelerator Facility), undated[a], “Jefferson Lab: Exploring the Nature of Matter,” <https://web.archive.org/web/20160730151051/https://www.jlab.org/>, accessed September 2016.

JLab, Hall D, undated[b], “The GlueX Experiment,” <http://web.archive.org/web/20161024055108/http://gluex.org/Gluex/Home.html>, accessed September 2016.

Kneller, J., and C. Volpe, 2010, “Turbulence Effects on Supernova Neutrinos,” *Phys. Rev. D*, Volume 82, 123004.

Krouppa, B., R. Ryblewski, and M. Strickland, 2015, “Bottomonia Suppression in 2.76 TeV Pb-Pb Collisions,” December 14, *Phys. Rev. C* 92, 061901(R).

LANL (Los Alamos National Laboratory), undated, “Los Alamos Neutron Science Center,” <http://web.archive.org/web/20161118010317/http://lansce.lanl.gov/>, accessed September 2016.

Lawrence Livermore National Laboratory, undated, “National Ignition Facility and Photon Science,” <http://web.archive.org/web/20160916033431/https://lasers.llnl.gov/>, accessed September 2016.

Lentz, E.J., S.W. Bruenn, W.R. Hix, A. Mezzacappa, O.E. Bronson Messer, E. Endeve, J.M. Blondin, J.A. Harris, P. Marronetti, and K.N. Yakunin, 2015, “Three-Dimensional Core-Collapse Supernova Simulated Using a $15 M_{\odot}$ Progenitor,” *Astro. J. Lett.*, Volume 807, Number 2, L31.

LIGO Scientific Collaboration, undated, “Introduction to LIGO & Gravitational Waves: Inspiral Gravitational Waves,” <http://web.archive.org/web/20161024024052/http://www.ligo.org/science/GW-Inspiral.php>, accessed January 2017.

Lovato, A., S. Gandolfi, J. Carlson, S.C. Pieper, and R. Schiavilla, 2016, “Electromagnetic Response of ^{12}C : A First-principles Calculation,” *Phys. Rev. Lett.* 117, 082501.

Matsui, T., and H. Satz, 1986, “ J/ψ Suppression by Quark-Gluon Plasma Formation,” October, *Phys. Lett. B*, Volume 178, 416.

McDonnell, J.D., N. Schunck, D. Higdon, J. Sarich, S.M. Wild, and W. Nazarewicz, 2015, “Uncertainty Quantification for Nuclear Density Functional Theory and Information Content of New Measurements,” *Phys. Rev. Lett.* 114, 122501.

McLerran, L.D., and R. Venugopalan, 1994, “Computing Quark and Gluon Distribution Functions for Very Large Nuclei,” March 1, *Phys. Rev. D* 49, 2233.

Melson, T., H.-T. Janka, R. Bollig, F. Hanke, A. Marek, B. Mueller, 2015, “Neutrino-Driven Explosion of a 20 Solar-Mass Star in Three Dimensions Enabled by Strange-Quark Contributions to Neutrino-Nucleon Scattering,” *Astro. J. Lett.* 808, L42.

Mocsy, A., P. Petreczky, and M. Strickland, 2013, “Quarkonia in the Quark Gluon Plasma,” February, *Int. J. Mod. Phys. A* 28, 1340012.

Moll, R., C. Raskin, D. Kasen, and S.E. Woosley, 2014, “TYPE Ia Supernovae from Merging White Dwarfs. I. Prompt Detonations,” *ApJ*, 785, 105.

MSU (Michigan State University), undated, “Facility for Rare Isotope Beams,” <https://web.archive.org/web/20160903021107/http://frib.msu.edu/>, accessed September 2016.

National Academies Press, 2013, *Nuclear Physics: Exploring the Heart of Matter*, National Research Council Report, Paperback, ISBN: 978-0-309-26040-4, DOI: 10.17226/13438, <https://www.nap.edu/catalog/13438/nuclear-physics-exploring-the-heart-of-matter>, accessed October 2016.

Nature, 2010, 466(7303):159–284, <https://web.archive.org/web/20120928074435/http://www.nature.com/nature/journal/v466/n7303/>, accessed September 2016.

Nature.com, undated, http://www.nature.com/polopoly_fs/7.2605.1328009973!/image/Neutron_web.jpg_gen/derivatives/nature_homepage/Neutron_web.jpg, accessed January 2017.

Nazarewicz, W., 2016, “Challenges in Nuclear Structure Theory,” *J. Phys. G* 43, 044002.

NSAC (National Science Advisory Committee), 2014[a], *Report to the Nuclear Science Advisory Committee: Neutrinoless Double Beta Decay*, April 24, http://science.energy.gov/~media/np/nsac/pdf/docs/2014/NLDBD_Report_2014_Final.pdf, accessed September 2016.

NSAC, 2014[b], *Assessment of Workforce Development Needs in the Office of Nuclear Physics Research Disciplines*, Subcommittee on Workforce Development, July 18, https://science.energy.gov/~media/np/nsac/pdf/docs/2014/NSAC_workforce_jul-18-2014.pdf, accessed January 2017.

Ohio State University, undated “NUCLEI: Nuclear Computations Low-Energy Initiative,” <https://web.archive.org/web/20161018215119/http://nuclei.mps.ohio-state.edu/>, accessed September 2016.

OLCF (Oak Ridge Leadership Computing Facility), undated, “Introducing Titan: Advancing the Era of Accelerated Computing,” <https://web.archive.org/web/20161017234311/https://www.olcf.ornl.gov/titan/>, accessed September 2016.

ORNL (Oak Ridge National Laboratory), undated, “Neutron Sciences,” <https://web.archive.org/web/20130823124412/http://neutrons.ornl.gov/>, accessed September 2016.

Perlmutter, S., 2012, “Nobel Lecture: Measuring the Acceleration of the Cosmic Expansion Using Supernovae,” *Rev. of Mod. Phys.*, Volume 84, 1127.

Pratt, S., E. Sangaline, P. Sorensen, and H. Wang, 2015, “Constraining the Equation of State of Superhadronic Matter from Heavy-Ion Collisions,” May 19, *Phys. Rev. Lett.* 114, 202301.

Rapp, R., 2011, “Theory of Soft Electromagnetic Emission in Heavy-Ion Collisions,” October, *Acta Phys. Polon. B* 42, 2823.

Read, J.S., L. Baiotti, J.D.E. Creighton, J.L. Friedman, B. Giacomazzo, K. Kyutoku, C. Markakis, L. Rezzolla, M. Shibata, and K. Taniguchi, 2013, “Matter Effects on Binary Neutron Star Waveforms,” *Phys. Rev. D* 88, 044042.

Roberts, L.F., D. Kasen, W.H. Lee, and E. Ramirez-Ruiz, 2011, “Electromagnetic Transients Powered by Nuclear Decay in the Tidal Tails of Coalescing Compact Binaries,” *Astro. J. Lett.*, Volume 736, L21.

Rosswog, S., 2013, “The Dynamic Ejecta of Compact Object Mergers and Eccentric Collisions,” *Philosophical Transactions of the Royal Society A*, Volume 371, 20120272. DOI: 10.1098/rsta.2012.0272.

Ruiter, A.J., K. Belczynski, S.A. Sim, W. Hillebrandt, C.L. Fryer, M. Fink, and M. Kromer, 2011, “Delay Times and Rates for Type Ia Supernovae and Thermonuclear Explosions from Double-Detonation Sub-Chandrasekhar Mass Models,” *Monthly Notices of the Royal Astronomical Society*, Volume 417, pp. 408–419.

Sadhukhan, J., W. Nazarewicz, and N. Schunck, 2016, “Microscopic Modeling of Mass and Charge Distributions in the Spontaneous Fission of ^{240}Pu ,” *Phys. Rev. C* 93, 011304(R).

Shepherd, M.R., J.J. Dudek, and R.E. Mitchell, 2016, “Searching for the Rules that Govern Hadron Construction,” *Nature* 534:487–493. DOI 10.1038, <http://www.nature.com/nature/journal/v534/n7608/abs/nature18011.html>, accessed September 2016.

Strohmayer, T., 2004, “Future Probes of the Neutron Star Equation of State Using X-Ray Bursts,” *AIP Conference Proceedings*, Volume 714, 245. DOI: 10.1063/1.1781035.

Takiwaki, T., K. Kotake, and Y. Suwa, 2012, “Three-Dimensional Hydrodynamic Core-Collapse Supernova Simulations for an $11.2 M_{\odot}$ Star with Spectral Neutrino Transport,” *Astro. J. Lett.*, Volume 749, 98.

USQCD (U.S. Lattice Quantum Chromodynamics), undated, “Particle and Nuclear Physics,” <https://web.archive.org/web/20161204175455/http://www.usqcd.org/>, accessed September 2016.

Vlasenko, A., G.M. Fuller, and V. Cirigliano, 2014, “Neutrino Quantum Kinetics,” *Phys. Rev. D*, Volume 89, 105004.

Wanajo, S., Y. Sekiguchi, N. Nishimura, K. Kiuchi, K. Kyutoku, and M. Shibata, 2014, “Production of All the r -Process Nuclides in the Dynamical Ejecta of Neutron Star Mergers,” *Astro. J. Lett.*, Volume 789, L39.

Wilson, K., 1974, “Confinement of Quarks,” *Phys. Rev. D* 10(8):2445.

Zingale, M., C.M. Malone, A. Nonaka, A.S. Almgren, and J.B. Bell, 2015, “Comparisons of Two- and Three-Dimensional Convection in Type I X-ray Bursts,” *ApJ*, 807, 60.

6 ACRONYMS AND ABBREVIATIONS

$0\nu\beta\beta$	neutrinoless double-beta decay
1-D	one-dimensional
2-D	two-dimensional
3-D	three-dimensional
4-D	four-dimensional
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
AMG	adaptive multi-grid
AMO	atomic, molecular, optical
AMR	adaptive mesh refinement
API	application programming interfaces
ASIC	application-specific integrated circuit
ASCR	Advanced Scientific Computing Research
ATLAS	Argonne Tandem Linac Accelerator System
B	baryon number
BER	Biological and Environmental Science
BES	Basic Energy Sciences
BiCGStab	stabilized biconjugate gradients
BNL	Brookhaven National Laboratory
BSM	beyond the Standard Model
C/O	carbon/oxygen
CAAR	Center for Accelerated Application Readiness
CC	coupled clusters
CCSNe	core-collapse supernovae
CEBAF	Continuous Electron Beam Accelerator Facility
CERN	European Organization for Nuclear Research
CG	conjugate gradient
CGC	Color-Glass-Condensate
CI	configuration-interaction
CISE	Computer and Information Science and Engineering
CP	[combined] charge-conjugation and parity
CPU	central processing unit
CUORE	Cryogenic Underground Observatory for Rare Events
CVMFS	Cern VM file system

DD	domain decomposed
DFT	density functional theory
DOE	U.S. Department of Energy
DUNE	Deep Underground Neutrino Experiment
EBE	event-by-event
ECP	Exascale Computing Project
EDM	electric dipole moment
EFT	effective field theories
EIC	Electron Ion Collider
EoS	equation of state
Fe	iron
Fermilab	Fermi National Accelerator Laboratory
FES	Fusion Energy Sciences
FRIB	Facility for Rare Isotope Beams
FY	fiscal year
GPD	generalized parton distribution function
GPU	graphics processing unit
GR	general relativistic
H	hydrogen
He	helium
HEP	High-Energy Physics
HISQ	highly improved staggered quark
HMC	hybrid Monte Carlo
HPC	high-performance computing
HPCDE	high-performance computing and data ecosystems
HRS	High-Resolution Spectrograph
I/O	input/output
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
IP Glasma	impact-parameter-dependent Glasma
JLab	Thomas Jefferson National Accelerator Laboratory
KATRIN	Karlsruhe Tritium Neutrino Experiment

L	lepton number
LANL	Los Alamos National Laboratory
LCF	Leadership Computing Facility
LHC	Large Hadron Collider
LIGO	Laser Interferometer Gravitational-Wave Observatory
LLNL	Lawrence Livermore National Laboratory
MD	molecular dynamics
MSU	Michigan State University
NASA	National Aeronautics and Space Administration
nEDM	neutron electric dipole moment
NERSC	National Energy Research Scientific Computing Center
NESAP	NERSC Exascale Application Program
Ni	nickel
NIF	National Ignition Facility
NN	nucleon–nucleon
NNSA	National Nuclear Security Agency
NP	Nuclear Physics
NRC	National Research Council
NS	neutron star
NS&R	nuclear structure and reactions
NSAC	Nuclear Sciences Advisory Committee
NSCI	National Strategic Computing Initiative
NSCL	National Superconducting Cyclotron Laboratory
NSF	National Science Foundation
NSNS	double neutron star
NUCLEI	Nuclear Computations Low-Energy Initiative
NVRAM	non-volatile random-access memory
OLCF	Oak Ridge Leadership Computing Facility
O-Ne-Mg	oxygen-neon-magnesium
ORNL	Oak Ridge National Laboratory
OSG	Open Science Grid
pQCD	perturbative QCD
QCD	quantum chromodynamics
QCDSP	QCD on digital signal processors
QCDOC	QCD on a chip




QGP	quark-gluon plasma
QKE	quantum kinetic equation
QMC	quantum Monte Carlo
RHIC	Relativistic Heavy Ion Collider
RXTE	Rossi X-ray Timing Explorer
SC	Office of Science (DOE)
SciDAC	Scientific Discovery through Advanced Computing
Si	silicon
SNe Ia	Type Ia supernovae
SNS	Spallation Neutron Source
T	time-reversal
TeO ₂	tellurium dioxide
Theta ESP	Theta Early Science Program
TMD	transverse momentum dependent distribution function
UNEDF	Universal Nuclear Energy Density Functional
USQCD	U.S. Lattice Quantum Chromodynamics (Consortium)
VMC	virtual Monte Carlo
WD	white dwarf
XRB	x-ray burst
XSEDE	Extreme Science and Engineering Discovery Environment

Units of Measure


fm	femtometer(s)
GB	gigabyte(s)
GeV	giga electron volt(s)
keV	kilo electron volt(s)
kg	kilogram(s)
$M_{\odot}$	Solar mass(es)
MeV	million electron volt(s)
ms	millisecond(s)
PB	petabyte(s)
TB	terabyte(s)
TeV	tera electron volt(s)

NP

NUCLEAR PHYSICS



APPENDICES: MEETING MATERIALS



An Office of Science review sponsored jointly by
Advanced Scientific Computing Research and Nuclear Physics



APPENDIX A: NUCLEAR PHYSICS ORGANIZING COMMITTEE AND MEETING PARTICIPANTS

A.1 NUCLEAR PHYSICS ORGANIZING COMMITTEE

Joseph Carlson	Los Alamos National Laboratory (Co-Chair)
Martin J. Savage	Institute for Nuclear Theory, University of Washington (Co-Chair)
Richard Gerber	National Energy Research Scientific Computing Center
James Osborn	Argonne Leadership Computing Facility
Katherine Riley	Argonne Leadership Computing Facility
Tjerk Straatsma	Oak Ridge Leadership Computing Facility
Jack Wells	Oak Ridge Leadership Computing Facility

A.2 NUCLEAR PHYSICS SPONSORS AND REPRESENTATIVES

Ted Barnes	DOE/Office of Nuclear Physics
Carolyn Lauzon	DOE/Office of Advanced Scientific Computing Research

A.3 NUCLEAR PHYSICS MEETING PARTICIPANTS

Katerina Antypas	National Energy Research Scientific Computing Center
Deborah Bard	National Energy Research Scientific Computing Center
Ted Barnes	DOE Office of Nuclear Physics
Michael Bernhardt	Oak Ridge National Laboratory
Amber Boehnlein	Thomas Jefferson National Accelerator Facility
David Brown	Lawrence Berkeley National Laboratory
Alan Calder	Stony Brook University
Joseph Carlson	Los Alamos National Laboratory
Susan Coghlan	Argonne National Laboratory
Mario Cromaz	Lawrence Berkeley National Laboratory
David Dean	Oak Ridge National Laboratory
William Detmold	Massachusetts Institute of Technology
Jason Detwiler	University of Washington
Sudip Dosanjh	National Energy Research Scientific Computing Center/ Lawrence Berkeley National Laboratory/
Huaiyu Duan	University of New Mexico
Adrian Dumitru	Baruch College City University of New York
Robert Edwards	Thomas Jefferson National Accelerator Facility
Jonathan Engel	University of North Carolina
George Fai	DOE Office of Nuclear Physics

Chris Fryer	Los Alamos National Laboratory
George Fuller	University of California, San Diego
Stefano Gandolfi	Los Alamos National Laboratory
Richard Gerber	National Energy Research Scientific Computing Center
Lisa Gerhardt	National Energy Research Scientific Computing Center
David Goodwin	DOE/Office of Advanced Scientific Computing Research
Rajan Gupta	Los Alamos National Laboratory
James Hack	Oak Ridge National Laboratory
Timothy Hallman	DOE
John Harney	Oak Ridge National Laboratory
Graham Heyes	Thomas Jefferson National Accelerator Facility
William Raphael Hix	Oak Ridge National Laboratory
Gustav Jansen	Oak Ridge National Laboratory
Richard Jones	University of Connecticut
Balint Joo	Thomas Jefferson National Accelerator Facility
Olaf Kaczmarek	University of Bielefeld
Daniel Kasen	University of California Berkeley/Lawrence Berkley National Laboratory
Kerstin Kleese van Dam	Brookhaven National Laboratory
Douglas Kothe	Oak Ridge National Laboratory
Thorsten Kurth	National Energy Research Scientific Computing Center/ Lawrence Berkley National Laboratory
Jerome Lauret	Brookhaven National Laboratory
Carolyn Lauzon	DOE/Office of Advanced Scientific Computing Research
Glenn Lockwood	Lawrence Berkeley National Laboratory
Kathryn Mace	ESnet/Lawrence Berkeley National Laboratory
Paul Mantica	Facility for Rare Isotope Beams, Michigan State University
Bronson Messer	Oak Ridge National Laboratory
Paul Messina	Argonne National Laboratory
Anthony Mezzacappa	University of Tennessee/Oak Ridge National Laboratory
Bogdan Mihaila	National Science Foundation
Swagato Mukherjee	Brookhaven National Laboratory
Hai Ah Nam	Los Alamos National Laboratory
Witold Nazarewicz	Facility for Rare Isotope Beams/Michigan State University
Esmond Ng	Lawrence Berkeley National Laboratory
Konstantinos Orginos	College of William and Mary/Thomas Jefferson National Accelerator Facility
James Osborn	Argonne National Laboratory

Thomas Papatheodore	Oak Ridge Leadership Computing Facility
Suzanne Parete-Koon	Oak Ridge National Laboratory
Peter Petreczky	Brookhaven National Laboratory
Robinson Pino	DOE
Jeff Porter	Lawrence Berkley National Laboratory
Scott Pratt	Michigan State University
Ji Qiang	Lawrence Berkley National Laboratory
Sofia Quaglioni	Lawrence Livermore National Laboratory
David Richards	Thomas Jefferson National Accelerator Facility
Katherine Riley	Argonne National Laboratory
Gautam Rupak	Mississippi State University
Martin Savage	Institute for Nuclear Theory
Bjoern Schenke	Brookhaven National Laboratory
Rocco Schiavilla	Jefferson Lab/Old Dominion University
Soeren Schlichting	Brookhaven National Laboratory
Nicolas Schunck	Lawrence Livermore National Laboratory
Cory Snively	National Energy Research Scientific Computing Center
Jay Srinivasan	National Energy Research Scientific Computing Center/ Lawrence Berkley National Laboratory
Patrick Steinbrecher	Brookhaven National Laboratory
Charles Still	Lawrence Livermore National Laboratory
Tjerk Straatsma	Oak Ridge National Laboratory
Michael Strickland	Kent State University
Sreenivas Sukumar	Oak Ridge National Laboratory
Sergey Syritsyn	Thomas Jefferson National Accelerator Facility
James Vary	Iowa State University
Venkatram Vishwanath	Argonne National Laboratory
Jack Wells	Oak Ridge National Laboratory
Julia White	Oak Ridge National Laboratory
Stefan Wild	Argonne National Laboratory
Michael Zingale	Stony Brook University
Jason Zurawski	Lawrence Berkley National Laboratory/ESnet

APPENDIX B: NUCLEAR PHYSICS MEETING AGENDA

TUESDAY, JUNE 14

6:30 Dinner coordination meeting for organizing committee and breakout leads at hotel (optional, but recommended)

WEDNESDAY, JUNE 15

7:30 – 8:30 Registration, Refreshments

8:30 – 9:00 Welcome from DOE and the Program Committee

9:00 – 9:20 Purpose and Goals of the Meeting
Barbara Helland, ASCR

9:20 – 10:00 Nuclear Physics Science Drivers and Perspectives
Ted Barnes, DOE NP

10:00 – 10:30 Break

10:30 – 11:30 ASCR Computing Facilities Plans and Technology Roadmap (includes Q&A/discussion)
ASCR Facilities Roadmap and Technology Trends; What Requirements Reviews Can Address and Influence
Katherine Riley, ALCF

Exascale Computing Project (ECP) Update
Paul Messina, Argonne

11:30 – 12:00 DOE Data Management Policies and Requirements
Laura Biven, Senior Science and Technology Advisor, Office of the Deputy Director for Science Programs

12:00 – 1:00 Working Lunch

DOE Application Readiness and Portability
Tjerk Straatsma, OLCF

Charge to Working Groups
Richard Gerber, NERSC

1:00 – 3:30 Breakout Sessions

- Nuclear Astrophysics
George Fuller, University of California, San Diego; Bronson Messer, Oak Ridge National Laboratory
- Experiment and Data
Amber Boehnlein, Thomas Jefferson National Accelerator Facility; Jason Detwiler, University of Washington; Paul Mantica, Michigan State University; and Jeff Porter, Lawrence Berkeley National Laboratory
- Nuclear Structure
Joe Carlson, Los Alamos National Laboratory; Witek Nazarewicz, Michigan State

- Cold QCD
Robert Edwards, Thomas Jefferson National Accelerator Facility; Martin Savage, Institute for Nuclear Theory, University of Washington
- Hot QCD
Peter Petreczky, Brookhaven National Laboratory; Michael Strickland, Kent State

3:30 – 4:00	Break
4:00 – 5:00	Open discussion with ASCR Facilities Director Barbara Helland
5:00 – 6:00	Open discussion
6:00	Dinner on your own

THURSDAY, JUNE 16

8:00 – 8:30	Refreshments
8:30 – 10:00	Preliminary feedback from breakouts (15 minutes each, plenary)
10:00 – 10:30	Break
10:30 – 12:00	Cross-cutting Issues <i>Joe Carlson and Martin Savage</i>
12:00 – 1:00	Working Lunch
1:00 – 3:00	Breakout Sessions: Develop outlines and identify main breakout findings; incorporate into white papers
3:00 – 3:30	Break
3:30 – 4:45	Reports on Wednesday Breakouts – Breakout Leads (15 minutes each)
4:45 – 5:30	Integrate breakout findings into draft review findings and recommendations **End for Most Participants**

FRIDAY, JUNE 17

8:30 – 12:00	Instructions for co-chairs <i>Joe Carlson and Martin Savage</i> Co-chairs, leads continue to work on report Completed outline for each breakout and overall report Deadlines and assignments <i>Richard Coffey, ALCF</i>
--------------	---

APPENDIX C: NUCLEAR PHYSICS WHITE PAPERS

The following white papers were submitted by the authors listed below in advance of the Nuclear Physics Exascale Requirements Review to guide both the agenda and workshop discussions.

C.1 White Paper Addressing Nuclear Astrophysics

- C-3 Bronson Messer, Oak Ridge National Laboratory, Oak Ridge, Tennessee,
 and George Fuller, University of California, San Diego, California

C.2 White Paper Addressing Data/Experiments

- C-11 Amber Boehnlein, Thomas Jefferson National Accelerator Facility, Newport News,
 Virginia; Jason Detwiler, University of Washington, Seattle, Washington; Paul Mantica,
 Michigan State University, Lansing, Michigan; and Jeff Porter, Lawrence Berkeley
 National Laboratory, Berkeley, California

C.3 White Paper Addressing Nuclear Structure and Reactions

- C-17 Joseph Carlson, Los Alamos National Laboratory, Los Alamos, New Mexico; and
 Witek Nazarewicz, Michigan State University, Lansing, Michigan

C.4 White Paper Addressing Cold Quantum Chromodynamics

- C-24 Robert Edwards, Thomas Jefferson National Accelerator Facility, Newport News,
 Virginia; and Martin J. Savage, Institute for Nuclear Theory, Seattle, Washington

C.5 White Paper Addressing Hot Quantum Chromodynamics

- C-33 Peter Petreczky, Brookhaven National Laboratory, Upton, New York;
 and Michael Strickland, Kent State University, Kent, Ohio

C.1 White Paper Addressing Nuclear Astrophysics

Nuclear Physics White Paper on Nuclear Astrophysics

**Bronson Messer (Oak Ridge National Laboratory) and
George Fuller (University of California, San Diego)**

1. Science drivers for nuclear astrophysics

Observations of nucleosynthetic yields are perhaps the most common data used to infer the origin of the elements that compose the sun and the solar system. However, the formation and dissemination of these nuclei depend on the physics of astrophysical environments where they are produced. These include steady burning in stars, explosive burning in stellar explosions (for example, core-collapse supernovae, thermonuclear supernovae, X-ray bursts), neutron star mergers, and black hole formation.

On the basis of fundamental nuclear physics and observed cosmic abundances [2,3], it is known that roughly half of the isotopes heavier than iron were formed via a slow neutron-capture process (the s-process), with most of the remaining half due to a rapid neutron-capture process (the r-process) and a scattering of rarer isotopes ascribed to what was originally thought to be a proton-capture process (the p-process). We have considerable confidence that the s-process occurs within the stable burning shells of giant stars. In contrast, our incomplete understanding of the r- and p-processes represents a critical long-standing open problem in physics; addressing it requires new experimental measurements of highly radioactive isotopes coupled to advanced computational models of the potential astrophysical sites.

Several astrophysical environments are suspected to produce the conditions needed for r-process nucleosynthesis. The most promising scenarios are: (1) the collapse of oxygen/neon/magnesium or iron stellar cores with their associated supernovae; (2) the decompression of nuclear matter ejected during the violent merger of two neutron stars (or a neutron star and a black hole); and (3) the winds from gaseous disks accreting onto black holes formed in failed supernovae or neutron star mergers. Simulations of each of these scenarios with quantitative physical fidelity rely on similar toolkits: magnetohydrodynamics (MHD), thermonuclear kinetics, the equation of state (EoS) of nuclear matter, radiation transport of neutrinos and photons, and a general relativistic description of gravity. Current models that rely on approximations to all of this physics already stress petascale architectures. A comprehensive suite of new physics implementations is needed, designed to take maximal advantage of hybrid multi-core and many-core architectures.

Many explosive astrophysical phenomena are powered by runaway thermonuclear reactions occurring in material in the interior or on the surface of stars. These events include classical novae, type I X-ray bursts (XRBs), and type Ia supernovae (SNe Ia)—all involving matter transferred from a companion onto a compact object giving rise to a thermonuclear runaway. Modeling these events also requires multiphysics, multiscale hydrodynamics codes, where the additional physics (e.g., gravity and nuclear reaction networks) and the scales of these problems sometimes necessitate subgrid-scale models for unresolvable physics like flames and background turbulence.

Our ignorance of the properties of atomic nuclei near the drip lines has motivated a succession of Nuclear Science Advisory Committee (NSAC) Long Range Plans to call out the need to build an experimental facility capable of filling this void [1]. The Facility for Rare Isotope Beams (FRIB) is currently

under construction and scheduled to begin operation in 2022. Experiments at FRIB will be capable of probing many of the nuclei of interest in the r-process and p-process, but maximizing the return on the multimillion-dollar investment in FRIB requires focusing the experimental efforts on the species of greatest interest. From the astrophysical perspective, the foremost need is to establish the r-process path, an effort that requires establishing the r-process sites.

2. Science challenges expected to be solved in the 2020—2025 time frame using extant high-performance computing ecosystems

With Exascale computational resources, an r-process site census can be carried out (i.e. some sites can be substantially ruled out and the likely r-process path or paths can be established). The most likely sites are the neutrino-driven outflows formed in the aftermath of core-collapse supernovae (CCSNe) and neutron star mergers.

CCSNe mark the death of massive ($M > 10M_{\odot}$) stars, yielding bright and energetic explosions and the births of neutron stars (NSs) and black holes. After several million years of evolution, a massive star's core is composed of iron (and similar "iron-peak" elements) from which no further nuclear energy can be released by fission or fusion. Outside the iron core are shells representative of previous burning stages—a silicon shell, oxygen shell, etc., out to a helium shell surrounded by an envelope composed of hydrogen. At the base of the silicon shell, nuclear burning continues, growing the iron core below. When the mass of the iron core reaches the limiting Chandrasekhar mass, it starts to collapse. During the collapse, the inner core becomes opaque to neutrinos and surpasses the density of atomic nuclei (approximately $2.5 \times 10^{14} \text{ g cm}^{-3}$), reaching densities where individual nuclei merge together into nuclear matter. Above nuclear density, the nuclear EoS stiffens, and the core rebounds like an over-compressed spring, launching a shock from the newly formed proto-NS. The rebound shock progresses outward through the rest of the in-falling core, heating and dissociating the in-falling nuclei to free nucleons and radiating a large burst of neutrinos. Thermal energy removed from the shocked material by neutrinos and nuclear dissociation halts the progress of the shock rendering it a standing accretion shock (SAS) with a radius of 100–200 km about 50 ms after it was launched. Deposition of energy behind the shock by neutrinos emitted by the proto-neutron star, aided by multi-dimensional effects (for example, convection and the stationary accretion shock instability), has been invoked to "revive" the shock and continue its progress to the stellar surface. However, models computed until now have all made marked approximations to the fundamental physics. To reduce the employed computational resources, current supernova simulations adopt various, and in some cases crude, approximations to reduce the high dimensionality of Boltzmann neutrino transport. Approaches to general relativistic gravity are approximate, at best, except in 1D. Limits on spatial resolution have led to questions regarding the impact of under-resolved turbulent flows on the explosion mechanism. The physical fidelity in modeling all of these ingredients can be substantially improved at the Exascale.

In CCSNe, the characters of both the Stationary Accretion Shock Instability (SASI) and neutrino-driven convection in 2D CCSNe simulations are different than in 3D. The nature of the buoyant engine, including its viability, is unknown for symmetry-free 3D configurations across the initial mass function. The SASI is confined to the sloshing mode by axisymmetry, and turbulent convection develops unrealistically large and strong convective rolls, which may work to reinforce each other and result in more stable accretion streams and a strong buoyant engine. Recently, several groups have obtained realistic explosions in 3D simulations [8], but only a narrow slice of the full space of progenitor mass, metallicity, rotation rates, and initial magnetic fields has been explored to date. Exascale resources are

required to extend this coverage through ensembles of 3D simulations and for a systematic study of the details of nucleosynthesis in “normal” CCSNe.

The merger of two neutron stars (NSMs) is the currently favored model to produce short-duration gamma-ray bursts. In addition to being a potential source of r-process yields, NSMs are believed to be one of the main sources for gravitational waves [4]. With the first detection of gravitational waves, the potential of using joint gravitational and electromagnetic wave emission to constrain properties of dense nuclear matter has moved from a theoretical possibility to a real scientific program. To actually use these mergers as probes (and understand their potential for gamma-ray bursts and nucleosynthesis), detailed models of the merger process and the electromagnetic emission are essential. Although current models of neutron star mergers are becoming more sophisticated, these models make simplifying assumptions in the physics: either simplifying the effects of general relativity, the behavior of matter at nuclear densities (equation of state), or the neutrino transport (cross-sections or radiation transport). Exascale resources offer the computational power to relax all of these assumptions, just as the same physical phenomenon will be better modeled in CCSN simulations.

Exascale computation will also allow us to identify the likely progenitor systems for SNe Ia. SNe Ia are the thermonuclear explosion of a C/O white dwarf (WD) in a binary system. These bright explosions rival the luminosity of the host galaxy, and their use as distance indicators led to discovery of the accelerating expansion of the Universe. However, theoretically, there are fundamental uncertainties in our understanding of SNe Ia—in particular, what is the progenitor? A single WD accreting from a companion star (single degenerate scenario), or two WDs that in-spiral and merge or collide (the double degenerate scenario)? No conclusive observation of the progenitor system exists, but in all cases the majority of the carbon and oxygen in the WD(s) is converted into Fe/Ni and intermediate-mass elements like silicon, and this nuclear energy release unbinds the star. For many years, the Chandrasekhar-mass single degenerate model (henceforth called the Chandra model) saw the most attention. The Chandrasekhar mass is the maximum possible mass of a WD, so explosions at this mass imply SNe Ia would be alike. Contemporary observations suggest diversity, and it is not clear if nature makes SNe Ia this way—massive C/O WDs in binary systems are rare. Accordingly, future simulations will need to explore all explosion models.

The standard picture for mergers has the WDs in-spiraling as gravitational radiation removes orbital energy. The less massive star will become tidally disrupted, and the more massive star will accrete this material. A long-standing concern is that when this mass transfer begins, thermonuclear burning can ignite at the edge of the star, converting it to O/Ne/Mg, and leading to the collapse of the WD into a neutron star, instead of giving rise to a SNe Ia. This system is inherently 3D, and only through the kinds of detailed simulations possible with Exascale platforms can we understand the dynamics of the mass transfer and thereby assess the feasibility of this model. Sub-Chandra models have the advantage that systems with a moderate mass WD ($0.8\text{--}1.0 M_{\odot}$ [solar masses]) are abundant, and this model also can reproduce the known delay-time distributions and SNe Ia rates—something the Chandra model struggles with. In this model, the WD accretes a layer of helium on its surface that detonates, sending a compression wave into the underlying C/O WD and igniting a second detonation that unbinds the star. Understanding how the ignition takes place in these shells when a realistic 3D convective field is realized, as well as the character and subsequent evolution of that detonation, is a goal for the immediate future.

3. Science challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems

A quantitative theory of the r-process—including precise yields per event, definitive calculations of the neutron excess, and other specifics—will likely require the application of *sustained* Exascale resources. Such predictions will require the ability to perform detailed, fully coupled, multi-angle and multi-energy radiation hydrodynamics simulations with full GR and very large (i.e., containing thousands of nuclear species) nuclear networks. Magnetic fields are also important to both the CCSN and NSM problems, and although a fully resolved, full-system, MHD calculation is probably beyond what can be accomplished in the next decade, simulations should be able to advance our understanding of this physics considerably.

As may be evident in the previous discussion, neutrino physics plays an important role in the evolution of compact object systems, e.g., in both CCSNe and binary neutron star mergers. In fact, most of the energy released in these events (approximately 10% of the *rest mass* of these objects) is emitted as neutrinos. The way energy and entropy is transported and deposited in these systems is neutrino-flavor dependent. Moreover, the energy spectra and fluxes of the different neutrino flavors may differ, so that if one flavor of neutrino transforms into another flavor in the supernova environment, the heating and energy balance can be altered, and this may also influence the local ratio of neutrons to protons, a key quantity for nucleosynthesis, especially for the r-process. A vexing numerical issue is that neutrino-flavor transformation in these environments is fiercely nonlinear, with the flavor states of the neutrinos determining how they interact and evolve. This nonlinearity stems from the “self-coupling” of neutrinos, a misnomer referring to the fact that neutrinos can forward scatter from other neutrinos, providing a flavor-dependent potential driving flavor conversion.

Large-scale simulations of coherent (meaning forward scattering only) neutrino-flavor transformation have revealed the possibility of a rich zoo of collective neutrino-flavor transformation modes in some epochs of compact object evolution. These simulations involve solving 10^7 to 10^8 nonlinearly coupled Schrodinger-like equations for neutrino flavor. This is a mature subject at this point [5,6]. Though issues remain unresolved, e.g., spatial and temporal instability in flavor transformation, the frontier in this subject now is incorporating these calculations into more realistic supernova and binary neutron star environments so as to gauge the effects of transformation on hydrodynamics/heating, neutron-to-proton ratio and nucleosynthesis. These goals are achievable with existing computational capabilities over a five-year time span.

Another level of complexity in these calculations comes when considering the role of neutrino direction-changing scattering-induced de-coherence in neutrino flavor conversion. The full quantum kinetic equations (QKEs) governing the evolution of neutrino flavor and spin (left- to right-handed, meaning neutrino-to-antineutrino conversion for Majorana-type neutrinos) in dense media have been derived recently. These are effectively generalizations of the Boltzmann neutrino transport equations, but the QKEs encode flavor mixing and neutrino spin. Simple limits of these equations have already been treated numerically, with fascinating results. Even a modest fraction (e.g., 1 in 10^3) of neutrinos suffering a direction-changing scattering can modify flavor transformation, now sending quantum mechanical flavor information *inward*, toward the neutron star, and effectively morphing the otherwise coherent neutrino-flavor transformation problem from an initial value problem, with neutrino flavors specified on the neutrino sphere, into something more akin to a boundary value problem. This is the so-called neutrino “halo” effect. The halo has been successfully simulated for O-Ne-Mg core collapse supernovae. A five-year goal is to model the halo effect in Fe-core collapse supernova.

A full quantum kinetic treatment of neutrino-flavor transport integrated into the state-of-the-art supernova and NSM simulations is an aspirational target for sustained Exascale and beyond resources. However, QKE calculations are a frontier issue in modeling early universe physics at the weak decoupling epoch. The high degree of symmetry in this environment reduces the complexity of QKE calculations, allowing simulations to be performed with existing supercomputing capabilities. Progress on these spatially homogeneous and isotropic conditions may allow insights into how to take advantage of increased memory and speed capabilities of Exascale Ecosystems to effect exploration of QKE physics in compact objects.

4. The most important needs for computing and data in 2020 and 2025 relative to today

The large number of degrees of freedom required in nuclear astrophysics simulations leads to considerable requirements for memory capacity and bandwidth. Programming models that allow code architects to effectively place and move data to various levels of the memory hierarchy are of key importance, as task-based and other asynchronous algorithms are explored to handle the variety of multi-physics inherent in these simulations.

Portable programming models are desired, as all practitioners can and do make use of the full panoply of available platforms, and expect to continue to do so into the future (Table 1). Without portable approaches, code maintenance and testing regimes will become too complex to manage for the small groups that develop the requisite codes.

Table 1: Resource currently available to the nuclear astrophysics program and expected requests in 2020 and 2025

Category	Present	2020	2025
Conventional core hours (M)	100	400	400
Multi-core/GPU node hours (M)	20	2000	16000
Aggregate memory (TB)	60	600	6000
Data read/write per run (PB)	1	20	100
Maximum I/O bandwidth (TB/s)	0.1	1	4
I/O fraction	1-2%	1-2%	1-2%
Scratch file space (PB)	3	60	80
Permanent space (PB)	3	60	80
Archival storage (PB)	3	60	80

5. Data and its storage, transmission, (real-time) analysis, and processing

The sheer number of physical variables evolved in supernova and NSM simulations translates directly into large near-line data storage volume requirements. Convolved with the increased spatial resolution anticipated on Exascale platforms, these requirements become very large indeed. I/O bandwidth

requirements are not as stringent, as the volume of computation brought about by better physical fidelity will serve to keep the data rate to disk essentially constant.

Though the data to be shared with experimental, observational, and other communities is relatively small when compared to other kinds of astrophysical simulation (e.g., large-scale structure), improved methods for data transmission will be important to ensure dissemination.

6. What top *three* computing ecosystem aspects will accelerate or impede your progress in the next 5–10 years?

Accelerate	Why?
1. FLOP-rich architectures	Nuclear astrophysics still has unrealized parallelism to tap into most of our problems and formulations.
2. Algorithms	Some algorithms used in nuclear astrophysics have become <i>de rigueur</i> over time because of computing power constraints. We now have an opportunity to explore new algorithms, including ones that were set aside earlier (e.g., Monte Carlo transport) for want of computing power.
3. Workforce development	Ubiquitous parallelism for all programming tasks will lead to a demand for the kind of training provided by nuclear astrophysical simulation.

Impede	Why?
1. Immaturity/lack of consensus on programming models	Nuclear astrophysics codes must run on varied platforms to execute the full program of simulations required.
2. Lack of memory	Nuclear astrophysics problems typically involve multiple degrees of freedom per spatial grid point, while often also requiring high spatial resolution.
3. Workforce development	Without significant increases in funding, the U.S. strength in the field will be ceded to Europe, Japan, and others, making it an unattractive option for students.

7. Opportunities for collaboration between ASCR and Nuclear Physics

Algorithm development is a critical part of the process of modeling nuclear astrophysical explosions. For example, the tight coupling of reactions and hydrodynamics, especially in regimes of explosive nucleosynthesis, demands stronger coupling between these physics modules; otherwise, we risk realizing an unphysical thermodynamic state. The common prescription for coupling these processes is Strang splitting, taking advantage of the disparate time scales associated with hydrodynamics and


reactions. However, this disparity of time scales is not always so large in some explosive environments (e.g., in some stages of SNe Ia). New approaches beyond Strang splitting could be useful in these cases.

Because much of the insight that can be gleaned from the large data volumes produced in nuclear astrophysics simulations is dependent upon visualization of those data, new techniques and tools to accomplish new visual analyses will be crucial. Both implementations of existing techniques that are scalable and responsive, as well as new techniques for analysis are needed.


8. References

- [1] A. Aprahamian, et al. "Reaching for the horizon: The 2015 long range plan for nuclear science." INSPIRE-1398831 (2015).
- [2] E. M. Burbidge, G. R. Burbidge, W. A. Fowler, and F. Hoyle (1957), "Synthesis of the Elements in Stars," *Rev. Mod. Phys.*, 29 547–650.
- [3] A. G. W. Cameron (1957), Nuclear Reactions in Stars and Nucleogenesis, *PASP*, 69 201–222, Atomic Energy of Canada, Ltd., CRL-41.
- [4] LIGO Caltech, "Sources and Types of Gravitational Waves," <https://www.ligo.caltech.edu/page/gw-sources>.
- [5] H. Duan, G. M. Fuller, and Y.-Z. Qian, *Ann. Rev. Nucl. Part. Sci.* 60 , 569 (2010).
- [6] H. Duan, G. M. Fuller, and J. Carlson, *Computational Science and Discovery*, 1, 015007 (2008).


9. Images



The specific entropy of the stellar material from the 3D CCSN simulation of Lentz et al. 2015 *ApJ* **807** L31.



A volume rendering of radial velocities for some representative models of helium burning on the surface of white dwarfs. (Jacobs et al. arXiv:1507.06696 [astro-ph.HE])



Chemical structure of the ejected debris 100 s after ignition for a subset of SNe Ia explosion models with different ignition and detonation conditions. (D. Kasen et al. *Nature* **460**, 869–872 (2009) doi:10.1038/nature08256)

C.2 White Paper Addressing Experiments and Data

Nuclear Physics White Paper on Experiment and Data

**Amber Boehnlein (Thomas Jefferson National Accelerator Facility),
Jason Detwiler (University of Washington),
Paul Mantica (Michigan State University), and
Jeff Porter (Lawrence Berkeley National Laboratory)**

1. Science drivers for nuclear experiment and data

The scientific drivers associated with Experiment and Data were elaborated in the 2015 Long Range Plan for Nuclear Science and include the unfolding of the quark and gluon structure of hadrons and nuclei; realizing a predictive model of nuclei and their role in the cosmos; testing the particle–antiparticle nature of neutrinos and other fundamental symmetries and neutrino research that open new doors to physics beyond the Standard Model; and studying the properties and phases of quark and gluon matter in the high temperatures of the early universe. DOE has invested heavily in these science areas by providing world-leading research facilities, such as ATLAS, CEBAF, Relativistic Heavy-Ion Collider (RHIC), and the Fundamental Neutron Physics Beam Line and by planning for next-generation facilities in Facility for Rare Isotope Beams (FRIB), Electron-Ion Collider (EIC), and a ton-scale neutrinoless beta ($0\nu\beta\beta$) decay experiment. The successful science programs at these facilities critically depend on the optimization of experimental infrastructure. High-performance computing (HPC) plays an important role in the reliable and efficient execution of the nuclear physics experimental program as well as the rapid and well-grounded treatment of the resulting data. Areas of Experiment and Data that are impacted by HPC include reliable and safe accelerator operations, advanced detector development and simulation, and the analysis, sharing, and archiving of large data sets.

The science opportunities in Experiment and Data span the full science program, including nuclear structure, cold and hot QCD, nuclear astrophysics, and fundamental symmetries. Therefore, there is strong synergy between experiment and theory in all of these areas. Experiment and Data within nuclear physics shares crosscuts with HEP and BES, particularly in the design and optimization of complex particle accelerator systems. The nuclear physics and HEP communities also have a common need to realize efficient particle-detector arrays. The need to apply high-performance networks and near-detector HPC resources is critical for rapid feedback on event structure and detector systems performance.

2. Science challenges expected to be addressed in the 2020–2025 time frame using extant high-performance computing ecosystems

The science drivers in Experiment and Data cover the full breadth of scientific study supported by the nuclear physics research program. Even in this diverse environment, there are a variety of common methods and tools needed by nuclear physics experiments to acquire, process, and analyze future data from DOE nuclear physics facilities. Within the 2020–2025 time frame, HPC has the potential to impact the means and methods by which nuclear physics researchers collect, analyze, and manage data.

Regarding data collection, there is strong interest within the nuclear physics community today to engage in collaborative research into high-precision, real-time analysis of streamed data at DOE experimental facilities. Within the next four years, the goal is to realize some common tool sets for using HPC-like facilities at experimental sites to perform high-ingest data analysis and validation tasks on triggerless data streams. By 2025, the challenges presented by data streaming and validation may be well-in-hand with the implementation of end-to-end software frameworks that implement best practices for integrating data transport, messaging, compute services, and visualization components.

Detector simulations and data analysis are generally carried out in the nuclear physics community using the well-developed software packages Geant and ROOT. These packages, however, are not taking advantage of HPC architectures. At this time, it is critical for the community to initiate collaborative research efforts to adapt Geant4 simulation and ROOT-based data analysis frameworks to the coarse-grained parallelism on HPC systems. In the next few years, it should be possible to begin to realize vectorized detector simulations (Geant V) on HPC and models or best practices for data analysis on emerging architectures. By 2025, new software frameworks for fine-grained parallelism for detector simulations and data analysis will allow nuclear physics researchers to approach complex detector simulations and the processing of large data sets taking full advantage of the computational power afforded by HPC.

Data management challenges within the nuclear physics community center mostly on complex work flows and inherent limitations on network Infrastructure. The pathway to address these challenges starts with collaborative research with nuclear physics scientists into adaptive, fault-tolerant data-processing work flows that optimize resource utilization. Within four years, the community will realize a common-use work-flow tool set that connects experimental site operations, institutional clusters, and ASCR-supported facilities. By 2025, end-to-end software frameworks will be developed to operate at different data-processing stages to efficiently use available computing resources and adapt to real-time environments.

3. Science challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems

The basic currency in nuclear physics experiments is data, both from direct measurement and detector simulations. New insights are obtained by making significant improvements in the acquisition, generation and analysis of data to increase the statistical precision and improve the accuracy of the measurements. Experiments are being designed today that will support data rates in the next 5-10 years beyond the TB/sec level (at the LHC/RHIC, FRIB, JLab), which, for today's near-line streaming capabilities and offline processing capacities, present significant limitations. Those two aspects frame the science challenges that cannot be solved by extant computing ecosystems for the field in the 2020–2025 time frame.

The current extant system of global fitting for deriving parton distribution functions (PDFs) is being extended to the extraction of generalized parton-distribution functions (GPDs) and transverse-momentum dependent parton-distribution functions (TMD), collectively called Nucleon Tomography. Exploring the computational requirements to realize this program has only recently begun; however, it is clear that a simple extension of the old methodology will be inadequate.

4. The most important needs for computing and data in 2020 and 2025 relative to today

The most significant computing need in nuclear physics experiment will be for effective real-time processing capability to manage the large data rates expected in the 2020–2025 time frame. Rates at over 1PB/day will become the norm, outpacing the long-term storage capacities that are predicted to exist for the community. Furthermore, at those interaction rates, conventional trigger systems used today to select the most interesting events for storing and later analysis will, in many cases, not be able to adequately disentangle those events from the rest of the data stream. The nuclear physics experiment community recognizes this challenge, and effort is being directed to develop such near-line capabilities for triggerless data streams. However, the community is simultaneously engaged in ongoing experiments that make it difficult to do the research and development work needed for such a large paradigm shift in online operations. The community would benefit greatly from the development of better algorithms (such as Machine Learning methods) and data-processing tools for lossless real-time data reduction near the beam line.

As can be seen in Table 1, nuclear physics experiments do not require large amounts of overall computational capabilities to reach our scientific goals. At the same time, the data-processing capabilities required by the community are illustrated by the 10s to 100s PB of online storage capacity needed for managing access to the very large analysis-ready data sets. The community expects that the new methods and tools developed for processing online data streams can also be used to enhance the throughput in offline data analysis. Use of such methods will require an HPC ecosystem that supports network-friendly adaptive work flows needed in such large data-processing tasks.

Table 1: Resource currently available in the nuclear physics experimental program and expected requests in 2020 and 2025

	Current Usage	Future Usage: 2020	Future Usage: 2025
Computational core hours (conventional)	870 M	1.7 B	3.8 B
Computational node hours (homogeneous many-core)	10 M	140 M	270 M
Computational node hours (with GPU or accelerator)	22 M	86 M	170 M
Memory per node	128 GB	128 GB	1024 GB
Aggregate memory	300 TB	3.0 PB	5.0 PB
Data read and written per run	1.0 PB	5.0 PB	15 PB
Maximum I/O bandwidth needed	290 GB/sec	700 GB/sec	5.6 TB/sec
Percentage of runtime for I/O	50%	50%	50%
Scratch file system space needed	430 TB	1.0 PB	3.3 PB
Permanent online data storage	14 PB	64 PB	200 PB
Archival data storage needed	65 PB	240 PB	550 PB

5. Data needs including storage, transmission, (real-time) analysis, and processing

Experimental data usually takes the form of recorded or simulated detector signals, parameters describing or derived from those signals, and other associated metadata, including detector configuration information, systems-monitoring data, and Monte Carlo truth. The current and expected cumulative data volumes and maximum velocities for the nuclear physics experimental program are given in Table 1. The numbers are dominated by the needs of large, many-channel, high-rate systems, which pose challenges at every stage in data collection and data processing.

During data collection and signal recording, data velocity is typically so high that fast triggering schemes are engineered to decide when to record data and when to discard it. As detector systems grow in size, data become increasingly complex, and rates increase further, the discrimination of potentially overlapping signals from noise or background will only become more challenging. A combination of HPC and storage coupled to the detector by a low-latency, high-bandwidth network would enable a new data-streaming paradigm that would circumvent many of these challenges. The data would persist long enough to perform online validation, filtering, and data reduction to ensure high-quality data-taking at acceptable rates to permanent storage, enabling much higher luminosity experiments than are currently feasible.

Data must be moved from the experimental site to a computing infrastructure for offline processing. Robust wide-area-network connectivity and capacity (10–100s Gbps) are needed to handle the high data throughput as well as, for instance, the required distribution of software and environments, all the way to node-level processing. Work-flow management also poses a significant challenge, as work flows spanning security domains are possible only with the existence of and reliance on secured interfaces, federated identities, and common authentication protocols across sites, not only to access computing resources but also to appropriately share data among collaborators and other authorized researchers. The Open Science Grid (OSG) has addressed many of those challenges, with active network-monitoring collaborations with ESNNet and Internet2, software distribution with OASIS (CVMFS), and federated identities using X509 certification now under the CILogon CA. However, at present, OSG is not generally available at larger ASCR facilities.

Once on disk, the data must then be further calibrated, reprocessed, filtered, and subjected to statistical analysis. In addition, large processed data sets must be archived in high-performance storage systems, from which they can be periodically restaged to disk for reprocessing. Moreover, developing a full understanding of the detector response requires immense simulation efforts that often generate large output intended to emulate the recorded data stream and sometimes exceed the size of the experimental data itself. While the natural division of data into temporal “events” enables (in most cases) embarrassing parallel data processing and simulation that make efficient use of HPC resources, the challenge posed is typically in the required I/O to move data from storage to memory, and to write computed parameters and simulation results back to storage again. The introduction of parallelization for data mining on large data sets can alleviate some of these challenges.

6. The top three things in an HPC ecosystem in 5–10 years that are (1) required to facilitate research and (2) expected to impede progress

Facilitate	Why?
1. Use of efficient, portable programming models for HPC-like architectures	HPC-like architectures will be more prevalent in specialized compute facilities used for online data-taking operations. As a result, nuclear physics experiment frameworks and code bases should become more portable, making HPC systems more usable.
2. Emergence and growth of data awareness at HPC centers, providing many PB- to EB-scale online data repositories	The ratio of CPU capacity to hot online storage capacity is shifting in HPC centers toward greater data and I/O awareness. Such shifts make these resources more valuable to nuclear physics experiment use cases.
3. Network as a service to support new processing paradigms such as triggerless data streams and adaptive work flows	The growth in data awareness and real-time analysis of data streams can change the way networks are used in HPC systems for better support of data-processing tasks commonly found in nuclear physics experiments.

Impede	Why?
1. Lack of open, resilient high-bandwidth network capacity across various scales, including from the worker node all the way to the wide-area network	Dynamic, data-intensive work flows need to communicate and move data between various resources. HPC systems (and centers) generally restrict communication and data movement, making those resources not usable for nuclear physics experiment work flows.
2. Lack of acceptance by HPC centers for distributed trust systems that provide federated identities for seamless access to resources across security domains	Nuclear physics experiments establish collaborations that connect and provide access to multiple compute resources for distributed data-processing work. Although trust requirements at HPC centers are often satisfied by the experiment membership policies, ASCR facilities generally do not accept group authorization, reducing the usability of HPC centers in such work flows.
3. Workforce development and systems for knowledge sharing	Nuclear physics experiments often work independently, re-inventing the same sets of tools. This mode of work will not be manageable as the computing landscape becomes more complex with widespread use of HPC-like resources, which will require a better prepared workforce in computing science techniques and a system for sharing information between experiments.

7. Opportunities for collaboration between ASCR and Nuclear Physics

In many cases, there is a similar set of opportunities for experimental and observational data that appears in multiple domains, and those are good candidates for broad collaboration. Examples include the triggerless data stream mentioned in the previous section. To enable the effective use and deployment of petascale machines at the beam line will require tools and infrastructure that are able to provide for a seamless transition from local resources to the large HPC centers. At a fundamental level, having joint partnerships in developing Exascale capable algorithms for reconstructing and analyzing data would be extremely productive. We are entering an era where there is a match between the

number of detection channels and the parallelism of computing platforms. Hit-finding and -tracking algorithms could be completely rethought in this paradigm, which will be freed from a serial track-by-track approach with considerable built-in assumptions in order to keep the algorithm speed matched to the available resources. While faster reconstruction would be one goal, it might be more important to increase the track-finding efficiency, to increase the precision of the momentum measurements, and to be able to find tracks down to lower momentum.

Meeting the physics goals articulated above is going to require a tighter coupling of data analysis, often over multiple sources (both experimental and simulated) with theoretical models to perform global fits of the data. Currently, such use cases are handled in an ad hoc way and still include extracting data from published plots. It would be possible to articulate many possible areas for deep collaboration in computer science, applied mathematics, data organization, manipulation and analytics, multi-dimensional visualization, and intelligent work flows that capture parameter state and code versions of trial runs and enable “publication” of the final results. To elaborate on one topic, nuclear physics experimental data tend to be organized and formatted in ways that are convenient for data collection and storage technology, and this mindset influences offline data organization into large files and drives the development of the tools that are used to process data in files, often with external databases needed for ancillary information. This rigid data structuring is not well matched to Exascale computing nor especially well matched to modern data analytics. At the same time, there has to be a very significant improvement in user experience to motivate change. Experimentalists who do algorithm and infrastructure development often consider usability and user interface issues as window dressing that they don’t have time for. This generally leads to relatively few active analysts within collaborations. An excellent joint project would be defining and tackling the set of projects that move data reconstruction and analysis from the current serial mode executed with aging and difficult-to-use tools to data models and tools that are appropriate for modern computing architectures and modern visualization techniques and that enable a broader range of scientists to do analysis.

An additional set of projects could focus on the integration and data mining of sensor data for accelerator operations, comparable to the work that is being undertaken in the commercial sector.

8. Relevant Reference Material

A large number of case studies were collected from the Nuclear Physics Experiment community to serve as the basis for this white paper as well as the other documents and presentations that have come out of this review process. This was the first time that the experimental community was involved, and given the huge scope of not only the science represented by this subpanel but also (and more importantly) the varying computational techniques and requirements of the different experiments themselves, this white paper was written to be as inclusive as possible. It should be mentioned that not all communities found the template applicable and that not all requested Case Studies were provided. In the end, we relied on the subpanel’s expertise to identify the highest priority issues and needs from the collected material. Those Case Studies that the subpanel felt best addressed these issues were selected and edited for inclusion in the final report. For completeness, the full set of Case Studies, as-collected (with light editing for readability), will be posted indefinitely at this URL: <http://tinyurl.com/z3np73s>.

C.3 White Paper Addressing Nuclear Structure and Reactions

Nuclear Physics White Paper on Nuclear Structure and Reactions

Joseph Carlson (Los Alamos National Laboratory) and
Witek Nazarewicz (Michigan State University)

1. Science drivers for nuclear structure and reactions

The structure and reactions of atomic nuclei are at the forefront of many of the most important Department of Energy (DOE) experimental facilities coming online in the next decade. The Facility for Rare Isotope Beams (FRIB) at Michigan State University is under construction now and will probe the structure and reactions of very-neutron-rich nuclei. The goals of this facility are to improve understanding of the structure of atomic nuclei far from the valley of stability, to better understand how their properties impact the creation of the heavy elements in the universe in the r-process, and to use these nuclei for studies of fundamental symmetries and societal applications.

Next-generation double-beta ($0\nu\beta\beta$) decay experiments will probe the Majorana nature of neutrinos (whether neutrinos are their own anti-particle) to unprecedented precision. Observation of neutrinoless $0\nu\beta\beta$ decay, which is possible only for certain nuclei, will confirm the Majorana nature of neutrinos. The rate of the decay is extremely dependent upon the relevant weak-interaction nuclear matrix elements; thus theoretical input is critical to guide the construction of the experiment and to determine its final sensitivity. In addition, a successful observation coupled with an accurate determination of the matrix element would determine the neutrino's absolute mass scale in a laboratory experiment.

In addition, next-generation neutrino oscillation experiments are designed to determine the phases of the neutrino mixing matrix and the neutrino mass hierarchy. The Charge-Parity (CP- or equivalently, time-reversal) violating matrix element(s) have not yet been observed. The determination of this phase requires an accurate understanding of both neutrino- and anti-neutrino-nucleus cross-sections including their energy dependence. Accurate theories of nuclear structure and reactions, particularly the electroweak response of nuclei, are a critical ingredient in reaching these goals.

Beyond even all the large-scale experiments and facilities described above, a comprehensive picture of nuclear structure and reactions is required across a range of disciplines. Neutron star properties, including mass and radius, are being probed with various astrophysical observations, soon to include observations of the gravitational waves produced in neutron star mergers. Properties of dense matter are critical here and also for core-collapse supernovae, two possible sites of r-process nucleosynthesis. Studies of dense matter must be directly connected to theories of the structure and reactions of atomic nuclei. Determining the properties of neutron star crusts requires calculations of the properties of very inhomogeneous and very neutron-rich matter with its rich pasta-like phases.

Nuclear physics is used to probe extreme environments, like the interior of stars and the plasmas produced at NIF, and often reaction cross-sections in regions that are difficult to probe experimentally due to the Coulomb barrier or the unstable nature of relevant targets are required. The development of precise theories of nuclear reactions is crucial to understanding a host of phenomena in these environments. Nuclear fission is critical to many applications of nuclear physics and also to understanding astrophysical scenarios like fission cycling in the r-process. Finally, nuclei are used as

probes of fundamental symmetries and physics beyond the standard model in processes like beta decay. Modest investments in nuclear theory and computations can pay tremendous dividends in enhancing the sensitivity and hence the impact of these experiments.

In summary, the physics of nuclear structure and reactions addresses overarching questions of modern science [0]. It plays a crucial role across some of the most important nuclear physics and high-energy-physics experiments under construction in the US and worldwide, and theory and computing play a crucial role in taking maximum advantage of these very-large-scale investments.

2. Science challenges expected to be addressed in the 2020–2025 time frame using extant high-performance computing ecosystems

The structure and reactions of neutron-rich nuclei, including many of those to be studied at FRIB, are critical to better understanding nuclear structure and reactions more generally, and to the formation of the heavy elements in r-process nucleosynthesis. Important progress has been achieved in the last few years in *ab-initio* studies of these nuclei with realistic two- and three- nucleon interactions. Because of the three-nucleon interaction and of the weak binding of nuclei, these studies become more challenging as the neutron drip line is approached. Their reach is expected to expand dramatically, providing crucial information on nuclear binding and shell effects and also on nuclear reaction rates. Increased computing capabilities are required to obtain more refined models of interactions and currents, and to employ these in large-scale many-body calculations to accurately predict the properties of many nuclei near the drip lines. These studies will evolve to address more complex phenomena including collective excitations and reactions on unstable nuclei. Reactions including neutron capture, electroweak excitations and capture, and beta decays will play a critical role in efforts to understand both the nuclear physics of neutron-rich nuclei and their role in the r-process.

Important progress in understanding beta decay and $0\nu\beta\beta$ decay will be made. The nucleon axial current coupling g_A is artificially reduced (or “quenched”) in simple models to reproduce experimental beta decay rates. This reduction is associated with incomplete models of nuclear interactions and currents and incomplete many-body calculations. The role of the weak coupling in $\beta\beta$ -decay experiments is particularly crucial as g_A enters to the 4th power in the rate. In this time period, with the dramatic increase in available computing power, a quantitative explanation of nuclear beta decay is expected, along with its associated impact upon $0\nu\beta\beta$ -decay experiments. Accurate complete matrix element calculations for ^{48}Ca , a somewhat simpler nucleus than the heavier germanium and xenon nuclei to be used in the ton-scale $0\nu\beta\beta$ decay experiments, will be accomplished.

Neutrino scattering on nuclei is another critical topic in nuclear structure and reactions. A quantitative understanding of neutrino scattering in both the low- (astrophysical) energy regime, and the quasi-elastic regime for intermediate mass nuclei is expected to emerge. This will again require realistic models of nuclear interactions and currents, and also reliable solutions for the response of many-body nuclear targets. Results on the carbon nucleus, the main constituent in present-day detectors, should be available in the next several years. In this time frame, initial results for argon, the main detector element in future neutrino experiments such as DUNE, are expected.

Our ability to calculate light-nucleus reactions from first principles, including some results in regimes not experimentally accessible, will be greatly enhanced. Various reactions important in big-bang nucleosynthesis, solar neutrinos, nuclear astrophysics, and nuclear reactions important to NIF will be

calculated. Nuclear fission is a critical ingredient in both applications and nuclear astrophysics; progress in this area has been, and will continue to be, dramatic.

The application of high-performance computing has revolutionized theories of complex nuclei by both optimizing the input and carrying out advanced applications. The scientific goals in this time frame are to make validated predictions of neutron-induced fission and beta-decay rates, and provide the microscopic nuclear physics input for r-process nuclei. This requires the development of more realistic energy functionals of spectroscopic quality, including finite-range and/or three-body nuclear forces.

Dramatic increases in our ability to predict properties of dense matter and related astrophysical environments are expected during this time period. In addition to studying dynamics in dense matter including neutrino propagation, nuclear density functional theory will be employed to study physics over a much broader range of length scales than can be accomplished any other way. This is particularly crucial, for example, in the very inhomogeneous neutron star crust.

3. Science challenges that cannot be solved in the 2020–2025 time frame using extant high-performance computing ecosystems

Even given the important progress possible in the next 5 to 10 years, critical challenges will remain. Extending the *ab-initio* and density-functional studies to fully describe complex coherent quantum phenomena will require both exascale resources and a sufficient investment in a workforce required to optimally exploit them. Accurate modeling of nuclear structure and reactions is critical to predict the structure of and reactions of neutron-rich nuclei at the edge of the line of stability. This more refined picture is also required to understand the r-process with sufficient fidelity to differentiate between possible r-process sites including supernovae and neutron star mergers.

A dramatic increase in our understanding of electroweak processes in nuclei is expected in the next 5 to 10 years. However, a determination of the $0\nu\beta\beta$ decay matrix elements for germanium and xenon with sufficient accuracy to extract the absolute neutrino mass scale and to test for additional new physics contributions to $0\nu\beta\beta$ decay will require resources beyond those available in the 2020–2025 time frame. Similarly, the extensions of neutrino cross-sections to more exclusive processes, or to the hadronic regime, including pion production and the impact of resonance production, will require further resources.

Not all reactions will be accessible in the 2020–2025 time frame. For example, a full *ab-initio* calculation of carbon-alpha capture to form oxygen, a critical cross-section for nuclear astrophysics, will require new resources. Other complex reactions including a comprehensive understanding of both spontaneous and induced fission with a determination of mass fragment and neutron distributions will remain challenging.

Significant challenges are expected to remain in understanding dynamics in nuclear astrophysical environments at finite temperature. Neutrino propagation in supernovae and neutron-star merger environments is crucial to a full understanding of supernovae dynamics and the origin of the heavy elements.

4. The most important needs for computing and data in 2020 and 2025 relative to today

The nuclear structure and reactions community involves a diverse set of applications, including algorithms like quantum Monte Carlo, the no-core Shell Model (NCSM, configuration interaction) and related reaction theory, coupled cluster, and density functional theory. While these algorithms have somewhat different needs, common themes exist.

The availability of significant capability and capacity computational resources remains critical to all of the applications in nuclear structure and reactions. All major applications are expected to move dominantly towards multi-core and/or GPU architectures over the next 5 years. This community has a very good track record and significant experience moving towards new architectures and taking advantage of the largest available computational platforms. As shown in Table 1, the needs for conventional core hours are estimated to be reduced from approximately 750M to 300M over the next 5 years and then stabilize. All the major applications are expected to evolve toward multi-core and/or GPU architectures, and the resource needs are estimated below in node hours (rather than core hours used for conventional architectures). This year, approximately 20M node hours are used, primarily by coupled-cluster on GPU architectures. However, the resource needs for all applications are estimated to expand to 1,500M node hours in 2020 and 11,000 node hours in 2025. All major algorithms and codes in this area face scientific challenges that require a vast increase in computing power.

Table 1: Resource currently available to the nuclear structure and reactions program and expected requests in 2020 and 2025

Category	Present	2020	2025
Conventional core-hours (M)	750	300	300
Multi-core/GPU node hours (M)	20	1500	11000
Aggregate Memory (TB)	750	5000	10000
Data read/write per run (PB)	0.1	0.5	2
Maximum I/O bandwidth (TB/sec)	0.015	0.05	0.6
I/O fraction (%)	5	5	5
Scratch file space (PB)	0.02	2	8
Permanent file space (PB)	0.1	3	10
Archival Storage (PB)	0.02	5	10

Certain algorithms, particularly those associated with large-scale linear algebra using configuration-interaction, and resonating group algorithms also require a large increase in aggregate memory. These algorithms project out low-lying and ground states using sparse linear algebra techniques. The increased memory requirements are associated with more complete treatment of three-nucleon interactions and their significant increase in memory requirements, and with attacking larger problems, both larger nuclei and the calculation of nuclear reactions, which are more challenging than bound-state calculations for the same size nucleus. An increase in aggregate memory requirements from 750 TB today to 5,000 TB in 2020 and 10,000 TB in 2025 is estimated. As future increases in aggregate memory

may be achieved by increasing the overall number of nodes rather than increasing the memory per core, plans to structure our codes to deal with these upcoming architectures are under way.

We rely heavily on many standard libraries for these large-scale machines, including MPI and OpenMP, dense and sparse linear algebra packages like BLAS and LAPACK, efficient Fortran compilers, and parallel I/O libraries. Libraries have been created to fill our specific needs regarding load balancing and memory management. Abstractions and tools for performance portability will become increasingly more important as computer architectures evolve. Visualization of the results will also become increasingly important.

5. Data and its storage, transmission, (real-time) analysis, and processing

The data requirements for storage and analysis/post-processing are also increasing greatly. Essentially all major algorithms are generating increasing amounts of data that should be stored for later analysis. Examples include calculations of nuclear reactions from *r*-matrix analyses of nuclear reactions, calculations of electron- and neutrino-scattering from nuclei and other response function calculations relevant to electroweak interactions with nuclei, and calculations of interaction and norm kernels for multi-reference DFT. Storing and analyzing results are also critical for optimizing nuclear interactions and currents efficiently.

The largest data sets to be stored long term include results of large-scale calculations, including basis set amplitudes (eigenvectors) and relevant matrix elements or, for the case of quantum Monte Carlo algorithms, samples of the solution with particle coordinates and spin/isospin amplitudes. In either case these data sets can become quite large for the challenging problems being addressed.

While the fraction of time used for I/O is expected to remain fairly constant at around 5%, a significant increase in needs for data storage is anticipated, including scratch space, permanent file space, and archival storage from about 0.05 PB each today to approximately 10 PB in 2025. This increased need is shared across all scientific problems and codes.

6. The top three things in an HPC ecosystem in 5–10 years that are (1) required to facilitate research and (2) expected to impede progress

Facilitate	Why?
1. Leadership-class computing resources	Sufficient access to leadership-class computing resources on multi-core/GPU architectures is critical to progress in the field, and especially to meet the scientific challenges presented by next-generation experimental facilities. It is critical to maintain and expand the presence of this area in leadership-class computing.
2. Algorithmic and software research	The software that is required to perform efficiently and portably on next-generation large-scale computers includes MPI and OpenMP, and efficient languages and libraries to deal with the GPU and multi-core architectures. Libraries dealing with linear algebra, task-based computing, and load balancing and optimization are critical. Efficient Fortran compilers that can deal with multiple architectures are required.
3. Storage requirements	Moving into the Exascale era, sufficient storage is required for intermediate results and later analysis on a variety of machines (capability and capacity), and network infrastructure to move data to the appropriate analysis sites. Scratch space is needed for immediate analysis, while permanent/archival space is required for longer term storage and later analysis.

Impede	Why?
1. Inadequate workforce	Lack of a sufficient workforce is a significant risk for progress in the field. The required workforce includes a mixture of applied mathematicians and computer scientists, physicists, and those with mixed duties. The separate funding models for math/CS and physicists do not fully meet the needs of the required teams. Sustainable funding for scientists in all fields is required to meet the challenges of the Exascale Era.
2. Inadequate memory	NCSM and reaction calculations require large aggregate memory footprints. Progress in these areas will be diminished if sufficient aggregate memory is unavailable.
3. Inadequate portability	Portability across different architectures is important for flexibility going forward as the different architectures mature. This directly affects the ability to collaborate across machines and algorithms, including internal and external collaborations across methods and applications.

7. Opportunities for collaboration between ASCR and Nuclear Physics

One very important challenge is for ASCR and the nuclear physics community to work together to develop a capable and diverse workforce. While graduate fellowships and other programs are very valuable, more work on creating new positions at the boundary of physics and computer science/applied math is needed. The two different offices naturally focus on their own areas of expertise, but maximum progress can be achieved only with excellent career paths spanning the two disciplines. These must be made attractive and be stable, as expertise in computing typically translates into many opportunities.

Another challenge is creating scalable and portable software directly relevant to our programs. These include, for example, task-based parallelism and load-balancing techniques, sparse and dense linear algebra, and optimization techniques. Many of these problems are common to different areas in the Office of Science and beyond. It is important to carefully identify these software capabilities required by all areas in nuclear physics. Partnerships between nuclear physics and ASCR could also work together on the level of programming models, exploring different models in terms of performance portability and scalability.

Finally, requirements and common techniques for software, data, and their storage and maintenance are very important. Storage needs of all types will grow rapidly in the next decade; some of the important data will be at experimental facilities. Efficient means of exchanging and moving data among facilities would be extremely valuable. Common resources and methods for data and code storage should enable provenance of data and reproducibility of results and also enable future efforts to build efficiently on past success.

8. References

- [0] *Nuclear Physics: Exploring the Heart of Matter*, National Research Council Report, National Academy Press, Washington, D.C., 2013.
- [1] The priorities for the Nuclear Physics program funded by DOE and NSF are outlined in the recent 2015 NSAC [Long Range Plan for Nuclear Science](http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf).
http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf.
- [2] 2014 High Energy Physics P5 Report. <http://www.usparticlephysics.org/p5/>.
- [3] Facility for Rare Isotope Beams. <http://www.frib.msu.edu/>.
- [4] FRIB: "Opening New Frontiers in Nuclear Science." http://frib.msu.edu/sites/default/files/FRIB_Opening_New_Frontiers_in_Nuclear_Science.pdf.
- [5] NSAC Report on Neutrinoless Double Beta Decay. http://science.energy.gov/~media/np/nsac/pdf/docs/2014/NLDBD_Report_2014_Final.pdf.
- [6] Fermilab Neutrino Program. <http://www.fnal.gov/pub/science/particle-physics/experiments/neutrinos.html>.
- [7] Deep Underground Neutrino Experiment. <http://lbnf.fnal.gov>.
- [8] The NUCLEI SciDAC 3 project. computingnuclei.org.
- [9] DOE SciDAC program. <http://www.scidac.gov>.

C.4 White Paper Addressing Cold Quantum Chromodynamics

Nuclear Physics White Paper on Cold QCD

Robert Edwards (Thomas Jefferson National Accelerator Facility) and
Martin J. Savage (Institute for Nuclear Theory)

1. Science drivers for cold QCD

Precision lattice QCD studies of the structure, spectra, and interactions of mesons, baryons, and nuclei, and of the fundamental symmetries of nature, are now crucial to the U.S. nuclear physics research program. Entering an era in which lattice QCD calculations will be routinely performed at the physical quark masses and include electromagnetism, the current program is closely tied to the U.S. experimental program in nuclear physics and to the scientific milestones established by the Nuclear Science Advisory Committee (NSAC), and provides crucial support to other areas of nuclear and high-energy theory. Over the next few years, the cold QCD program will evolve to provide support for the Facility for Rare Isotope Beams (FRIB), the planned Electron-Ion Collider (EIC), ton-scale $0\nu\beta\beta$, and other experiments that are the priorities in the field, as identified in the 2015 NSAC Long Range Plan [1]. Providing theoretical support is critical to capitalizing on investments made in the U.S. experimental program.

Lattice QCD calculations focused on hadron spectroscopy of mesons and baryons, including the study of exotic states, plays a pivotal role in developing our understanding of QCD. They align well with, and motivate aspects of, the 12-GeV program at Jefferson Lab, and are evolving to address the anticipated needs of a future EIC. Results from recent experiments have upended the conventional picture that hadrons have a simple classification scheme, which has ignited a firestorm of interest [2]. Lattice QCD calculations are critical to this effort, suggesting the existence of new types of particles, and address a unique challenge for our understanding of QCD: how does the complexity of the hadron spectrum emerge from the interactions of the quarks and gluons? A key goal is to identify the presence of “hybrid” mesons and baryons, exotic states of QCD in which the gluonic degrees of freedom are manifest, and whose discovery is the primary motivation for the GlueX experiment at Jefferson Lab.

A precise knowledge of the structure of the light hadrons is essential for many aspects of subatomic physics, from the electronic structure of atoms through the response of detectors to neutrinos, to the design of the next-generation hadron colliders. Such an understanding includes the complete and precise decomposition of the mass of the proton, the decomposition of its spin, and a complete three-dimensional tomography of the nucleon to be realized through the determination of generalized parton distribution functions (GPDs) and transverse-momentum dependent parton-distribution functions (TMDs). These studies support and complement the experimental programs at JLab, the Relativistic Heavy-Ion Collider (RHIC), and the Large Hadron Collider (LHC), and will be essential in mapping the gluonic structure of nucleons and nuclei in the Electron-Ion Collider (EIC) era. An understanding of the nature and strength of the interactions of the light hadrons with the carriers of the weak force is critical to the success of experiments focused on determining the properties of neutrinos. Related calculations are critical to searches for Dark Matter and other beyond the Standard Model (BSM) physics.

The forces that describe the stable and long-lived nuclei have been tightly constrained through decades of experiment. However, lattice QCD calculations are expected to play a critical role in future

refinements of the nuclear forces and interactions. Work on light nuclei, light hypernuclei, and nuclear forces is expected to complement and support the planned FRIB experimental program. Providing the fundamental input into nuclear many-body calculations through effective field theories and chiral nuclear forces, these calculations are essential in predicting the structure and decays of nuclei, nuclear reaction rates, the response of nuclei in electroweak processes, and the behavior of matter in extreme astrophysical environments, such as binary neutron-star mergers that are expected to be observed by the Laser Interferometer Gravitational-Wave Observatory (LIGO) [3].

The observed asymmetry between matter and antimatter requires the non-conservation of baryon number, B , and charge conjugation and parity combined (CP) during non-equilibrium dynamics in the earliest moments of the universe. The CP violation and $B+L$ violation present in the Standard Model (where L is lepton number) are insufficient to produce the observed asymmetry, and current and planned experiments are designed to further explore T , CP, B , and L violating processes. The new U.S.-based DUNE experiment, which is expected to start operation in 2024, will search for B -violating proton decay, and neutron-antineutron oscillation experiments are also envisaged. An experiment to measure the neutron Electric Dipole Moment (nEDM), a probe of time-reversal violation, is under development at the Spallation Neutron Source (SNS) at Oak Ridge. A new flagship initiative in nuclear physics is a ton-scale experiment to search for $0\nu\beta\beta$ of nuclei that, if observed, would provide explicit evidence of the violation of lepton number. To constrain the nature of the associated new physics, the nEDM and the rates of proton and $0\nu\beta\beta$ decays induced by the possible new operator structures must be calculated by using lattice QCD.

2. Science challenges expected to be addressed in the 2020–2025 time frame using extant high-performance computing ecosystems

A key objective of hadron spectroscopy is the determination of resonance properties of the hadron spectrum. The restriction of lattice QCD calculations to finite spatial volumes enables a mapping of the energy dependence of scattering amplitudes from which resonance properties can be determined. Predictions obtained from these calculations will guide experimental searches at JLab's GlueX experiment and confront future predictions in other sectors of QCD. In the charmonium sector, future calculations will address the origin of unexpected new experimentally observed resonances, requiring a new understanding of the emergent dynamical degrees of freedom from the quarks and gluons.

Hadron structure calculations will precisely determine fundamental quantities characterizing the nucleon, including form factors, moments of the parton densities, helicity, and transversity distributions, moments of GPDs and TMDs providing a three-dimensional map of the nucleon. They will support and complement the experimental programs at JLab 12 GeV, RHIC-spin, and Fermilab. With the JLab and RHIC experimental programs evolving into an EIC program, the focus of research will increasingly emphasize the gluonic structure of the nucleon and nuclei, and calculations of gluonic structure are expected to play a crucial role in EIC design and in the subsequent analysis of the experimental data. The current experimental discrepancy between different measurements of the proton charge radius (see, for example, Ref. [4]) will be resolved by lattice QCD calculations.

The spectra, properties, and interactions of light nuclei, multi-nucleon and neutron-rich systems, as well as hypernuclei, will be calculated at and near the physical quark masses, including electromagnetism, leading to a significant refinement of the nuclear forces through the use of effective field theories. The results of these studies, through collaboration with nuclear many-body theorists, will be incorporated

into nuclear structure and reaction calculations of quantities of importance in nuclear physics to support and complement the experimental nuclear and high-energy physics programs planned for JLab, FRIB, the EIC, and LBNF/DUNE and in searches for the $0\nu\beta\beta$ of nuclei. They will tightly constrain the three-neutron and four-neutron interactions and provide precision hyperon-nucleon forces, quantities that are vital in predicting the behavior of dense matter. The responses of nuclei to electroweak probes are key quantities in the analysis of neutrino detection, and the multi-body aspects of these interactions will be precisely determined. The same will be true for the interactions of Dark Matter candidates.

The thrust of the fundamental symmetries effort is to provide theoretical support to the experimental program aimed at identifying and exposing the sources of physics beyond the standard model. CP violation originating from BSM physics manifests itself at low energies through local CP-odd operators, and their impact on the structure and interactions of the nucleon and light nuclei will be calculated. In most BSM scenarios, the violation of baryon number induces proton decay, and the new DUNE experiment will be particularly sensitive to the decay mode $p \rightarrow \nu K$. To support the ton-scale $0\nu\beta\beta$ experimental program, lattice QCD calculations in few nucleon systems that provide crucial inputs into nuclear many-body calculations of $0\nu\beta\beta$ nuclear decay rates will be performed.

Common to these research thrusts are the ensembles of gluon-field configurations. There is significant coordination between these calculations in the design and use of the configurations. The sampling of the gluons is a crucial component of all of the anticipated calculations, and the use of the computational resources made available to the field is closely coordinated in order to optimize scientific output. It is expected that this coordination will continue through the Exascale era.

3. Science challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems

After 2025, calculations of the structure of hadrons will enter the high-precision era, approximately coinciding with the expected turn-on of the experimental program at an EIC. The valence-quark observables will be calculated with extremely high precision where required. For instance, g_A and its associated form factor will become calculable at the one-part-in-a-thousand level, a precision required to refine calculations of the pp-fusion rate that dominates solar burning. It is expected that calculations supporting and complementing the EIC program will be required during the operation of the EIC, both to further reduce the uncertainty in existing quantities and to make first predictions for more complex quantities. In the area of hadron spectroscopy, the recently discovered unconventional exotic-meson states are characterized by unnaturally small energies. While the next 5–10 years is expected to see remarkable progress in deciphering their nature and identifying what is expected to be towers of partner states, high-precision calculations of these systems will require resources beyond 2025. In the area of the physics of nuclei and nuclear forces, it is the lightest nuclei and systems of several nucleons that will be calculated prior to 2025. Electroweak interactions in these systems will be calculated at the percent level. Further refinement of the nuclear forces, particularly the high-precision multi-neutron forces, including electromagnetism, and precision electroweak studies will require resources extending beyond 2025. These are required for first-principles calculations of the structure and behavior of cold dense matter at the densities generated in core-collapse supernova and neutron stars. If the $0\nu\beta\beta$ decay of nuclei or nEDM is observed, a next generation of detectors will focus on precisely constraining the structure of the interactions. It will be necessary to launch a second generation of lattice QCD calculations in larger nuclear systems and with a more extensive operator structures to facilitate the refinement and identification of the underlying BSM physics.

4. The most important needs for computing and data in 2020 and 2025 relative to today

Lattice QCD calculations for nuclear physics proceed in three fairly distinct phases. The first is the generation of a representative ensemble of gauge configurations, requiring capability computational resources. The second is their analysis, which involves the computation of quark propagators, requiring the linear system solution of the Dirac-matrix with a large number of right-hand sides. The third step is the contraction of these solution vectors. The resource requirements of the propagators and contractions are comparable to those of the gauge generation, but can use capability and capacity resources.

Gauge generation is a critical component of all the nuclear physics campaigns and requires Leadership-Class capability resources. It is strong-scaling limited, and as we proceed into the Exascale era, the software must be both performance-portable and architecturally aware in order to exploit the available hardware. The key data types are fields at the sites and links of a regular hypercubic grid, partitioned into equal-sized domains per node. A key component of both the gauge generation and propagator analysis phases is the solution of the Dirac-equation—a large sparse linear with a dimension proportional to the lattice volume. A key challenge in the gauge generation phase is strong-scaling these solvers. Several methods have been developed to mitigate this challenge, relying on domain decomposition-based pre-conditioners and other methods of communication avoidance. Research into these methods will continue into the Exascale era. To enable performance portability between architectures, a method has been developed based on just-in-time (JIT) compilation which allows data layouts to be tuned to the architecture where the code is to be run, and uses the LLVM compiler framework to generate code for both NVIDIA GPUs and multi- and many-core CPUs at runtimes using the appropriate back-end code generator of LLVM.

The program envisaged over the period 2020–2025 will require the order of billions of node hours measured in units of Titan GPU hours. In the 2020 time frame, we envision producing 1,000 configurations separated by 10 trajectories for electromagnetically neutral lattices, not including isospin breaking, with spatial extents of 6 fm, 8 fm, and 10 fm, and lattice spacings of 0.077 fm and 0.06 fm. In the time frame between 2020 and 2025, additional neutral ensembles with lattice spacing 0.045 fm will be produced, and production of electromagnetically charged ensembles at 0.077 fm will begin. In the 2025 time frame, electromagnetically charged ensembles with lattice sizes of 6 fm and 8 fm at a lattice spacing of 0.06 fm will be generated, and generation of 10 fm charged ensembles at this lattice spacing will begin. The requirement for gauge generation in 2020 is estimated as 1.125 B Titan GPU hour equivalents and 4.3 B Titan GPU hours in 2025. Analysis is dominated by the time to solve the Dirac-equation for a large number of right-hand sides. This is a challenging computational task that requires the development of architecturally optimized Dirac solvers supporting nearly 1M right-hand sides. The third phase involves the contraction of the solution vectors and requires a carefully coordinated work flow to mitigate I/O demands. Intermediate paging requires investigations into new database and I/O technologies. The work flow in these last phases constitutes an in-situ-like analysis. Future algorithm development efforts will focus on different evaluation strategies that are more optimal for many-body baryon operator constructions. Overall, the combined analysis cost is roughly 3 times the gauge generation cost shown in Table 1, giving a total compute cost of approximately 4.5 B Titan GPU node hours in 2020 and 17.2 B Titan GPU hours in 2025. Computational and storage requirements are presented in the table below. There will be a continued need for 3rd party libraries for linear algebra (BLAS) routines, eigenvectors, and efficient FFT evaluation.

Table 1: Resource currently available to the cold QCD program and expected requests in 2020 and 2025

Code: gauge generation, analysis/propagators, contractions	Current Usage	Future Usage: 2020	Future Usage: 2025
Computational core hours (conventional)	350 M	300 M	300 M
Computational node hours (homogeneous many-core)	Reported in GPU	Reported in GPU	Reported in GPU
Computational node hours (with GPU or accelerator, using Titan GPU Node hour equivalents)	20 M	4.5 B	17.2 B
Memory per node	128 GB	128* GB	128* GB
Aggregate memory	8 TB	250 TB	400 TB
Data read and written per run	3 TB	250 TB	800 TB
Maximum I/O bandwidth needed	1 GB/sec	100 GB/sec	100 GB/sec
Percentage of runtime for I/O	10%	20%	20%
Scratch file system space needed	10TB	2 PB	8 PB
Permanent online data storage	500 TB	8.7 PB	31 PB
Archival data storage needed	1 PB	62 PB	315 PB

* Assumes that local volume is unchanged.

5. Data needs including storage, transmission, (real-time) analysis, and processing

The data generated by lattice QCD simulations fits into three distinct categories with different life cycles, access, and analysis patterns. The primary data generated are the gauge field configurations, which are used for all subsequent analyses. The primary need here is for long-term archival and curation based on the simulation parameters. Typically, generated gauge-field data remains active for 5–10 years and is stored both at the producing institution and at sites of the lattice QCD hardware project. There are presently efforts at curation, such as the International Lattice Data Grid effort and the Gauge Connection archive at NERSC. In terms of use in analysis, primarily read-only access is needed. In large throughput-oriented analysis projects, it may be that several of the analyses are simultaneously using a given set of data. The secondary data generated are the propagators. This data is typically too large to save, especially for the configurations expected in the 2020–2025 time frame. Hence, it is transient and survives until it is used to construct correlation functions. Primarily, the need here is for temporary storage until the correlation function construction phase is completed. The third and most important data are the resulting correlation functions. These tend to be small, compared to the propagators, and different for each particular analysis. Typically, to determine a physical quantity one will need a large

number of correlation functions, and the data-handling problem becomes one of data selection (query), cleaning (subset selection, binning), and statistical analysis. The challenge here is to have the data in such a format to enable this manipulation. Naive file-based approaches do not always scale, and current approaches use database technologies for efficiency. In the 2020–2025 time frame, it may be that better performance could be obtained using database services, rather than files.

As a large amount of data is expected to be generated and used quite broadly, data tracking, curation, and genealogy are becoming important issues. The ability to track data from its creation through its use and copying, to its destruction will be important to have been implemented before entering the Exascale era. Further, ensuring the data is secure is essential.

6. The top three things in an HPC ecosystem in 5–10 years that are (1) required to facilitate research and (2) expected to impede progress

Facilitate	Why?
1. Capacity- level analysis resources are essential moving into the Exascale era	The cost of analysis, which can use capacity or capability resources, is comparable to that of gauge generation, which can only be carried out on capability resources. Several such capacity-level resources will be needed, both at the Exascale facilities and institutions, including US QCD and universities. Local computational resources at PI institutions and those operated by US QCD at JLab, BNL and Fermilab have proven invaluable for the rapid turnaround and testing of new ideas and algorithms, circumventing the need for the preparation and approval of computing resource proposals (whose turnaround time is typically months). Such machines continue to play a crucial role in training junior scientists and attracting them to the field.
2. Performance-portable programming models, and algorithmic and software research	Exascale-relevant software development of architecturally aware, performance-portable, software infrastructure is needed. The software will be designed in a layer structure supporting high-level constructs written over a lower level API implemented in architecturally specific, optimized codes. The lower levels will rely on hybrid parallelism, with communications support and OpenMP/ACC/CUDA implementations. 3 rd party libraries such as linear algebra routines and continued support for advanced features proposed for C++ will also be required.
3. Support for efficient I/O	Moving into the Exascale era, analysis campaigns and contraction steps require significant coordination of compute and I/O–paging to disk. Workflow software should coordinate construction and caching temporaries to lower computational cost and I/O footprint.

Impede	Why?
1. Inadequate access to Leadership-Class compute resources	Lack of sufficient compute resources restricts the generation of gauge-field ensembles and limits physics campaigns. Gauge generation requires approximately 10 B GPU/KNL-equivalent node hours (measured in current Titan node hours), through the end of 2025, with a “burn rate” of between 1 B and 4 B GPU hours per year between 2020 and 2025.
2. Inadequate workforce	Lattice QCD calculations are workforce intensive, requiring HPC-trained teams of scientists with diverse skill sets. Postdoctoral fellows, laboratory scientists, and research faculty at universities, dedicated to the execution of lattice QCD calculations, are presently supported by the SciDAC program [5] and are expected to be additionally supported under the newly awarded Exascale Computational Project. Support of this nature is essential through the Exascale era.
3. Low I/O bandwidth compared to compute capabilities	I/O bandwidth will become more of an issue in the Exascale era as the analysis campaigns become more sophisticated.

7. Opportunities for collaboration between ASCR and Nuclear Physics

From the inception of the SciDAC program, collaborations among domain scientists in the area of cold QCD and ASCR’s SciDAC Institutes have been essential for the development of the program. Moving into the Exascale era, the need to strengthen these partnerships becomes paramount. In the Exascale era, one challenge is the development of scalable algorithms for such systems, especially in the case of multi-level algorithms like adaptive multi-grid methods. Typically these methods coarsen the original grid, resulting in potential strong-scaling challenges at the coarsest levels. One area of collaboration with ASCR would be to develop scalable linear solvers, which can scale well onto future Exascale systems, but which can match or exceed the performance of existing adaptive multi-grid methods. Another challenge of using Exascale systems is to maintain portability in general and performance portability in particular. Partnerships with ASCR could explore existing and envisaged programming models for performance portability both in terms of fine-grained threading and vectorization and in terms of exploiting upcoming memory hierarchies. The lessons learned could be beneficial to other areas facing these challenges. Even with performance-portable code, certain algorithms may suit particular hardware configurations. For example, one can imagine that in a domain-decomposed solver certain algorithmic parameters may be varied based on the architecture and the constraints of the memory hierarchy. One potential area of partnership would be the development and implementation of solver algorithms whose algorithmic parameters allow architecture-aware tuning. Finally, an area that will gain increasing importance is improving the I/O performance of our codes, especially in the propagator computation and correlation function construction phases. ASCR and ASCR facilities have considerable expertise with both efficient I/O libraries and data management. A partnership between the nuclear physics community and ASCR could evaluate and implement better I/O in our codes, explore the efficient use of burst buffers and could assist in investigations of using more modern data analysis tools (e.g., databases and database services) to analyze our correlation functions.

8. References

- [1] The priorities for the Nuclear Physics program funded by DOE and NSF are outlined in the recent 2015 NSAC “[Long Range Plan for Nuclear Science](#).”
http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf.
- [2] Searching for the rules that govern hadron construction, M.R. Shepherd, J.J. Dudek, and R.E. Mitchell, *Nature* 534, 487–493 (23 June 2016) doi:10.1038/nature18011.
- [3] LIGO Caltech, “Sources and Types of Gravitational Waves,” <https://www.ligo.caltech.edu/page/gw-sources>.
- [4] Scientific American, “The Proton Radius Puzzle,” <http://www.scientificamerican.com/article/the-proton-radius-puzzle/>.
- [5] DOE Office of Science, SciDAC, <http://www.scidac.gov>, and U.S. Lattice Quantum Chromodynamics, <http://usqcd.fnal.gov/index.html>.

9. Images




Figure 1. Contours of a correlation function of the Lambda baryon superposed on the fluctuations in a quenched action density. Image reproduced with permission from Martin Savage (INT).




Figure 2. A cartoon of nucleon-nucleon scattering. Image reproduced with permission from Martin Savage (INT).




Figure 3. A cartoon showing the naive structure of a conventional meson (left) and of a hybrid meson explicitly involving excitations of the glue. Image reproduced with permission from Robert Edwards (JLab).




Figure 4. Generating gluon fields: a cartoon of a "trajectory" around the equilibrium probability distribution. Image reproduced with permission from Balint Joo (JLab).

C.5 White Paper Addressing Hot Quantum Chromodynamics

Nuclear Physics White Paper on Hot QCD Requirements in the Exascale Era

Peter Petreczky (Brookhaven National Laboratory) and
Michael Strickland (Kent State University)

1. Science drivers for hot QCD

Using first-principles lattice QCD calculations, it is now well-established that strongly interacting matter undergoes a transition to a new state at a temperature of approximately 155 MeV at zero net baryon density ($\mu_B = 0$) and that this transition is a crossover rather than a phase transition [1, 2]. Furthermore, continuum-extrapolated calculations of the QCD equation of state (EoS) with physical quark masses (which required lattice QCD calculations performed over multiple years on the nation's fastest supercomputers) are available at zero net baryon density [3]. Such calculations have provided first-principles, non-perturbative information about the quark-gluon plasma (QGP) that is critically needed by the relativistic heavy-ion community.

The ongoing ultrarelativistic heavy-ion collision (URHIC) experiments at the Relativistic Heavy-Ion Collider (RHIC) and Large Hadron Collider (LHC) have provided a wealth of data that is being analyzed to determine and refine the properties of QGP. In particular, through comparisons of the predictions of relativistic viscous hydrodynamics for collective flow with experimental data, the heavy-ion community has inferred that the QGP generated in URHICs is nearly a perfect fluid, with a shear viscosity-to-entropy density ratio (η/s) that is close to the lower bound allowed by quantum mechanics. At the moment, a main focus of the heavy-ion program is to quantify the properties of the QGP through detailed comparisons between theoretical models and experimental measurements. Such comparisons permit the extraction of the properties of the QGP, such as transport coefficients from a variety of collision systems and collision energies. This program has been successful in describing data from the highest collision energy URHICs, which create a QGP with a small net baryon density.

Some of the important open questions being pursued by the hot QCD community are:

- What is the QCD phase diagram and the EoS for nonvanishing net baryon density?
- What are the transport coefficients of strongly interacting matter as a function of temperature and net baryon density?
- How are the early-time non-equilibrium dynamics of the QGP, and its subsequent approach to thermal equilibrium, realistically modeled?
- How can the deconfined nature of the new state of matter be established experimentally? At which energy scale can the QGP be understood in terms of weakly interacting quark and gluon quasiparticles?
- Are there hadronic excitations in the deconfined phase, and if so, what are their properties?

These important questions overlap with key scientific questions that drive the field of heavy-ion physics, as identified in the 2015 NSAC Long Range Plan [4].

In order to describe the collective flow experimental data, the hydrodynamic evolution has to start at early times, i.e. at 0.5–2 fm/c. Understanding how local thermal equilibrium can be achieved on such

short time scales, and what happens during the pre-equilibrium stages, is an outstanding theoretical challenge. It is known that, for the high-energy heavy-ion collisions, the initial state of the incoming nuclei can be described in terms of gluon saturation, which is characterized by an energy scale called the gluon saturation scale $Q_s \geq 1$ GeV [6, 7]. The large value of the gluon saturation scale (larger than the intrinsic scale of QCD, $\Lambda_{\text{QCD}} \sim 300$ MeV) implies that the early-time dynamics and perhaps the thermalization can be understood in terms of weak coupling methods. In the weak coupling picture, the gluon occupation number during the early stages is large, $f_g \sim 1/\alpha_s$, implying that the classical approximation can be used. Therefore, the process of thermalization can be studied using classical statistical simulations of the gauge theory (see, e.g., Ref. [8] for a recent work and references therein). Such simulations elucidate the approach to thermal equilibrium, degree of pressure anisotropy, and other such properties. However, current numerical simulations do not include the coupling of gluons to quark-antiquark pairs in the vacuum. Understanding how quarks are produced at the early stages of these collisions is an outstanding question, since about 2/3 of the entropy of thermalized QGP is contained in quark degrees of freedom. Therefore, fermions have to be included in the classical simulations. The inclusion of dynamical fermions in the classical statistical simulations is also important for understanding the electromagnetic response of the medium at early times and the anomalous transport, which is needed for understanding the chiral magnetic effect in RHIC heavy-ion collisions [9].

If the systems produced in heavy-ion collisions were in local thermodynamic equilibrium, the EoS and the transition temperature could be calculated by using lattice QCD. As mentioned previously, the EoS and the transition temperature are known for vanishing net baryon density. The next significant challenge is to extend these results to nonvanishing baryon density using the method of Taylor expansion, as direct lattice QCD simulations at non-zero baryon density are plagued by the infamous sign problem. The Taylor expansion method is also useful in exploring the phase diagram of QCD. The QCD transition at net-zero baryon density is a crossover, but this may change when the net baryon density or, equivalently, the baryon chemical potential, μ_B , is increased. Based on model calculations, it was conjectured that at sufficiently high value of μ_B , the transition changes from a crossover to a first-order phase transition, implying the existence of a critical point on the QCD diagram, where the first-order phase transition terminates and the transition is a second-order phase transition (see Fig. 1). In principle, it is possible that the transition remains a crossover up to very high value of μ_B ; that is to say that there is no phase transition as μ_B is increased (along the horizontal, μ_B , axis in Fig. 1) [10]. In this case, there is no critical end point in the QCD phase diagram. Lattice QCD calculations of the Taylor expansion coefficients could provide some evidence for the existence of the critical point or rule it out in a region of μ_B . The convergence of the Taylor series is limited by the nearest singularity in the complex μ_B plane. If the singularity is on the real axis, i.e. there is a critical point for some $\mu_B = \mu_{B,\text{crit}}$, the expansion coefficients are all positive for $T < T_c$. From the consequent orders of the Taylor expansion, one can provide estimates of the convergence radius, which can also be compared to the expectations based on a hadron gas. The Taylor expansion coefficients obtained so far are consistent with a hadron gas model within uncertainties. Therefore, there is no indication for the existence of a critical point from present lattice calculations.

Finding the critical point in the QCD phase diagram is the goal of phase II of the beam energy scan program (BES) at RHIC (BES II), which will take place in 2019–2020. By varying the energy of heavy-ion collisions, the net baryon density of the system can be changed and the region where the transition becomes first order may be explored experimentally. In the first phase of the BES program, which has recently concluded (BES I), some interesting non-monotonic behavior in several quantities were found as function of the beam energy. This might indicate the existence of a critical point.




Figure 1: A sketch illustrating the experimental and theoretical exploration of the QCD phase diagram. Although the experiments at highest energies and smallest baryon chemical potential achievable in the laboratory are known to cross from a QGP phase to a hadron gas phase through a smooth crossover, lower energy collisions can access higher baryon chemical potentials where a first-order phase transition line might exist [5].

The experimental uncertainties in BES II will be significantly smaller, and definite results confirming or disproving the non-monotonic behavior will become available. Linking such non-monotonic behavior to the existence of the critical point will require further theoretical investigations. Therefore, the above lattice QCD studies are critical to the scientific mission of RHIC.

The study of non-Gaussian fluctuations of conserved charges is an important tool for experimental searches for the critical point in the BES II program. These quantities can be calculated with lattice QCD using a Taylor expansion, as outlined previously. Some lattice QCD calculations of the non-Gaussian fluctuations have already been performed on configurations with a relatively large (coarse) lattice spacing and at a relatively low order in μ_B . In order to arrive at definite conclusions, these calculations should be extended to higher orders in μ_B , with further calculations on finer lattices to obtain the continuum limit. The interplay between the experimental program and the lattice QCD calculations at non-zero μ_B will be very important in the exploration of the QCD phase diagram.

To understand the properties of matter produced in heavy-ion collisions at RHIC and LHC, the study of penetrating probes, such as thermal photons, dileptons [11], and quarkonium [12], is important. In particular, the suppression of quarkonium production in heavy-ion collisions relative to appropriately scaled production in proton–proton collisions is considered an indication of QGP formation [13]. One of the themes of the 2015 NSAC Long Range Plan is probing the QGP at different length scales by studying quarkonium production in heavy-ion collisions. To interpret the experimental findings, scientists will need to know the properties of bottomonium at non-zero temperature (e.g., its thermal width) and the dissolution temperature of different quarkonium states. These are encoded in the meson spectral functions that can be extracted from Euclidean time correlation functions calculable with lattice QCD. In general, hadron correlation functions are important for heavy-ion phenomenology. Vector meson correlation and spectral functions in the light-quark sector provide information on the thermal photon and dilepton rates as well as on the electric conductivity of quark gluon plasma. Many of the calculations

of the meson spectral functions performed so far have been in the quenched approximation, and including dynamical quarks in future calculations will be necessary.

Since the systems created in heavy-ion collisions rapidly expand and are driven out of thermal equilibrium, knowledge of the EoS and transition temperature does not provide a complete description. The expansion has to be modeled using relativistic viscous hydrodynamics and, at later stages, hadronic transport models. The biggest unknowns in the hydrodynamic modeling are the transport coefficients, namely, the shear and bulk viscosities. While the EoS is known or will be known from lattice QCD, very little is known about the values and the temperature dependence of the various transport coefficients. The shear viscosity is expected to be large at low temperatures (as estimated in hadron gas models), as well as at very high temperatures, and likely has a minimum around the transition temperature. One possibility to map out the temperature dependence of the shear and bulk viscosities is through detailed comparison of the experimental data and hydrodynamic modeling. To do this, scientists need to parameterize the unknown physics related to initial conditions, temperature dependence of the viscosities, freeze-out conditions, and other quantities by a few dozen parameters and perform a large number of calculations including different possible values of these parameters. Transport coefficients at non-zero baryon density can also be studied using the experimental data from the BES. This will increase the computational requirements for two reasons. There will be data for several beam energies that will have to be analyzed. In addition, since at low energies the approximation of boost invariance does not hold, simulations using 3+1-dimensional hydrodynamics instead of 2+1-dimensional hydrodynamics need to be performed.

As mentioned previously, two years of runtime at RHIC, 2019–2020, will be devoted to the BES. These experiments represent exceptional opportunities for exploring regions at high net baryon density, and describing this lower energy regime requires profound improvements in modeling. At these energies, the approximations applied at high energies, based on approximate boost-invariance, are no longer valid, necessitating full three-dimensional hydrodynamics, which increases the numerical cost by two orders of magnitude. The importance of fluctuations, both from the initial conditions and those related to thermal phenomena, are essential for studying matter in this regime where there is a potential phase transition. The fluctuations require increasing the number of initial configurations by two to three orders of magnitude. Finally, at lower energy the matter spends more time in the hadronic region, $T < 154$ MeV, and hadronic mean fields need to be incorporated into the microscopic transport used to model the final stage and breakup of the collision. Combined, these model enhancements increase the numerical costs, both in cycles and in storage, by at least four orders of magnitude.

Recently, state-of-the-art statistical treatments have been applied to a global comparison of RHIC and LHC data with models, effectively constraining a high-dimension parameter space. These analyses have provided dramatic proof that experimental measurements from heavy-ion collisions can be used to constrain the QCD EoS and transport coefficients [14]. These calculations require running the full model at a large number of different points in parameter space. At high energies, this could be accomplished at smaller computing facilities, but with the approximately 100 times increase in computational requirements (mentioned above), leadership-class facilities are essential to realizing progress in this area and to fully exploiting the scientific investment in the BES at RHIC.

2. Science challenges expected to be in the 2020–2025 time frame using extant high-performance computing ecosystems

One of the goals of hot QCD research is to perform controlled calculations of QCD thermodynamics at non-zero baryon density using the Taylor expansion method. This will require performing calculations at smaller lattice spacings than today's, for example, on lattices with temporal extent $N_t = 12$ and 16. In addition, calculations of the expansion coefficients beyond the 6th order may need to be performed. These calculations should enable scientists to obtain the EoS at non-zero net baryon density in the continuum limit that is needed for hydrodynamical modeling of the collisions. This modeling is required to interpret results obtained with the BES II, the study of non-Gaussian fluctuations of conserved charges that are important for the search of the critical point in BES II, and finally to enable stringent constraints to be placed on the location of the critical point over a range of baryon chemical potentials.

Another important goal is to calculate hadron spectral functions, including the effect of dynamical quarks. As discussed previously, this is needed for the interpretation of experimental results from penetrating probes in heavy-ion experiments at RHIC and LHC. Quantifying the properties of QGP using penetrating probes is presently hindered by lack of knowledge of spectral functions.

The numerical study of equilibration and quark production at early times in heavy-ion collisions will be performed with realistic conditions including dynamical fermions. The study of the electromagnetic response of the system created in heavy-ion collisions will be performed with an aim to understand the charge separation observed in heavy ion collisions at RHIC and to determine whether it can be understood in terms of the chiral magnetic effect.

Finally, a comprehensive modeling of heavy-ion collisions with realistic initial conditions and 3-dimensional viscous hydrodynamics at many collision energies will be performed with the aim of extracting QGP transport coefficients as functions of temperature and baryon density.

The computational challenges, algorithms, and work flows of the four main subtopics discussed above—namely, lattice QCD at non-zero baryon density, studies of hadronic spectral functions, classical statistical simulations, and hydrodynamic simulations—are discussed in detail in four accompanying Case Studies.

3. Science challenges that cannot be solved in the 2020–2025 time frame using extant computing ecosystems

The study of the QCD phase diagram will most likely not be completed by 2025. The Taylor expansion technique is known to fail at large values of μ_b , and other techniques will be needed to circumvent the associated sign problem. It is possible that new algorithms, such as complex Langevin dynamics or Lefschetz Thimbles, will be sufficiently developed to make large-scale computations of the QCD phase diagram practical. The present studies of the QCD phase diagram, as well as those planned to be performed in 2020–2025, utilize the staggered fermion discretization of the quark fields, which does not respect the full chiral symmetry of QCD. Since chiral symmetry is essential for the QCD phase diagram, lattice QCD studies with chiral fermions, such as Domain Wall Fermions, are required. These discretizations of the quark fields require at least 10 times the computational resources required for staggered quarks.

Calculations of transport coefficients with lattice QCD are extremely challenging and require new techniques that possibly will not be available in the 2020–2025 time frame. Such calculations, however, are extremely valuable for understanding the physics of heavy-ion collisions at RHIC and LHC.

4. The most important needs for computational and data in 2020 and 2025 relative to today

The computational needs in the area of hot QCD are currently pushing towards the Exascale. In Table 1, a summary of our current computational needs is shown, along with projections for 2020 and 2025. In what follows, the methodology that leads to the estimates shown in the table for various subprojects is explained.

Within the area of lattice calculations at non-zero chemical potentials and spectral functions, there has been a focused effort of code development and optimization for more than a decade now within the USQCD (U.S. Lattice Quantum Chromodynamics) consortium (which combines the high-energy physics and nuclear physics communities involved in lattice QCD calculations, <http://www.usqcd.org/>). Groups involved in the USQCD consortium have been successful in running on leadership-class facilities very efficiently, in particular, on many-core/accelerated architectures. The lattice QCD calculations at non-zero chemical potential are running on a Titan machine in OLCF and will use 7 M node hours this year. There is a need to perform these calculations at smaller lattice spacings, leading to a resource requirement increase by approximately a factor of 20 by 2020 and by approximately a factor of 320 by 2025. These anticipated resource needs drive the larger number of required node hours on many-core architectures shown in Table 1. The lattice QCD calculations of meson spectral functions, discussed in the corresponding case study, have been performed in the quenched approximation. It will be necessary to extend these calculations by including the effects of dynamical fermions, which will lead to increased resource needs of approximately a factor of 100 in 2020 and approximately 1,000 in 2025. Because of the very large lattices that are needed for these calculations, the corresponding computational resource requirements are large even for those implementing the quenched approximation and amount to about 150M BG/Q core hours. The only way to perform the calculations with dynamical fermions is to use multi-core architecture and further algorithmic optimizations. Using a standard conversion between BG/Q core hours and GPU node hours, 1.43 M Titan node hours are estimated to be required for the quenched calculations. Algorithmic improvements, such as the use of anisotropic lattices, could reduce the computational requirements by a factor of 5. Combining these requirements, lattice QCD calculations of the spectral functions are estimated to need 30 M Titan-equivalent node hours in 2020 and 300 M in 2025.

The classical statistical simulations of the early-time dynamics of a heavy-ion collision will require the inclusion of dynamical quark–antiquark vacuum fluctuations, leading to an increase of the required computational resources by 3–4 orders of magnitude (depending on the fermion formulation implemented). At present, these calculations only use conventional CPU cores. Rescaling current use by the corresponding factor results in the estimates for conventional architectures shown in Table 1. It is clear that the requirements obtained this way are prohibitively large. Note, however, that for classical statistical simulations there has been no focused effort for algorithmic and code optimizations. These could potentially lead to a speedup of more than a factor of 5. Furthermore, using multi-core nodes can reduce the costs. We assume that a multi-core node, such as a Titan node, is equivalent to 10 conventional cores.

Table 1: Resource currently available in the hot QCD program and expected requests in 2020 and 2025

	Current Usage	Future Usage: 2020	Future Usage: 2025
Computational core hours (conventional)	165 M	1600 M	31000 M
Computational node hours (homogeneous many-core)	7 M	170 M	2500 M
Memory per node	16 GB	128 GB	256 GB
Aggregate memory	70 TB	2 PB	5 PB
Data read and written per run	10 TB	1 PB	1 PB
Maximum I/O bandwidth needed	400 GB/sec	800 GB/sec	4 TB/sec
Percentage of runtime for I/O	<5%	<5%	<5%
Scratch file system space needed	500 TB	4 PB	40 PB
Permanent online data storage	300 TB	3 PB	24 PB
Archival data storage needed	100 TB	1 PB	4 PB

In the case of hydrodynamic modeling, the main driver for the estimated increase in computational requirements is the need to perform calculations at many values of trial (input) parameters in order to (phenomenologically) map the temperature dependence of the transport coefficients, as well as the need for full 3 dimensional calculations at lower center-of-mass energy. This will result in factors of 10 times and 100 times increases in computational needs by 2020 and 2025, respectively. While the computational costs for hydro modeling are an order of magnitude smaller than those for classical statistical simulations, they are still significant. Therefore, an effort in software optimization is required, as well as the porting to multi-core architectures. Assuming that this will lead to performance gains similar to those of classical statistical simulations, the total needs of hydrodynamic modeling and classical statistical simulations are estimated to be 31M node hours and 620M node hours in 2020 and 2025, respectively.

Therefore, the total computational needs for the field of hot QCD in 2020 and 2025 are estimated to be 200 M and 2800 M node hours in terms of Titan-equivalent node hours, respectively. Using the USQCD metrics, this translates into equivalent sustained performances of 6.2 Pflop years in 2020 and 86 Pflop years in 2025.

Many of the applications do not need an extreme number of cores, and therefore communications will not be a bottleneck. The memory per node will not be an issue, as larger grid sizes will run on larger partitions; only the aggregate memory will increase. I/O is not expected to be a major problem.

5. Data needs including storage, transmission, (real-time) analysis, and processing

There will be a significant increase in the amount of data that is required to be stored on scratch disks, as well as permanently, to perform the different types of calculations. A large part of this need is

generated by the lattice QCD calculations, where the gauge configurations need to be, first, generated and, second, stored to be used later in calculating different quantities such as the Taylor expansion coefficients and various correlation functions. The generated gauge configurations could be used for other studies and will be archived for long-term use.

6. The top three things in an HPC ecosystem in 5–10 years that are (1) required to facilitate research and (2) expected to impede progress

Facilitate	Why?
1. Capacity-level resources are essential moving into the Exascale era	Local computational resources at various institutions and those operated by USQCD at JLab, BNL and Fermilab have proven invaluable for the rapid turnaround and testing of new ideas and algorithms, circumventing the need for the preparation and approval of computing resource proposals (whose turnaround time is typically months). Such machines continue to play a crucial role in training junior scientists and attracting them to the field.
2. Performance-portable programming models, and algorithmic and software research	With the exception of the lattice QCD codes for calculations of Taylor expansion coefficients that can run on conventional CPU, KNL, and GPU, most hot QCD codes have been developed for conventional CPUs. To exploit various resources that will be available in the next decade, it will be essential to have codes that can be ported to different architectures and that run efficiently.
3. Adequate memory bandwidth	Currently the performance of many QCD codes is bandwidth bound. To take advantage of many-core architectures, adequate bandwidth will be needed.

Impede	Why?
1. Inadequate access to Leadership-Class compute resources	Lack of sufficient compute resources restricts the generation of gauge-field configurations as well as subsequent calculations that use these gauge configurations.
2. Inadequate workforce	HPC-trained teams of scientists with diverse skill sets are needed to develop performance-portable codes to fully explore the opportunities in the Exascale era. Currently such teams are not available.
3. Low memory and network bandwidth compared to on-chip compute capabilities	Most applications need to move large amounts of data from memory to compute nodes and/or across the network. The tendency is that the bandwidth cannot keep up with on-chip compute capabilities. The problem can only be partially mitigated by changing the programming model.

7. Opportunities for collaboration between ASCR and Nuclear Physics

In terms of readiness for many-core and GPU-accelerated systems, different components of the hot QCD calculations are at different stages. The lattice QCD calculations at non-zero baryon density have been using GPUs for many years and can also run on multi-core systems like Intel KNL thanks to dedicated efforts in BNL and in Bielefeld University. Lattice QCD calculations of the spectral functions have a lot in common with cold QCD calculations, and therefore the code and algorithm development will be closely coordinated with the cold QCD community.

However, classical statistical simulations of early-time dynamics and hydrodynamic simulations currently only run on conventional CPUs. At present, there are no concentrated efforts to port the codes to GPUs and other many-core architectures. However, initial efforts in this direction are under way. In order to continue these developments, dedicated staff will be necessary which could provide an opportunity for fruitful collaboration between ASCR and the nuclear physics community. In particular, ASCR could help in developing programming models for performance-portable codes to be used by researchers working on classical statistical simulations and hydrodynamic modeling.

8. References

- [1] A. Bazavov et al., *Phys. Rev. D* 85, 054503 (2012) [arXiv:1111.1710 [hep-lat]].
- [2] T. Bhattacharya et al., *Phys. Rev. Lett.* 113, no. 8, 082001 (2014) [arXiv:1402.5175 [hep-lat]].
- [3] A. Bazavov et al. [HotQCD Collaboration], *Phys. Rev. D* 90, 094503 (2014) [arXiv:1407.6387 [hep-lat]].
- [4] Nuclear Science Advisory Committee (NSAC), "Reaching for the Horizon," The 2015 Long Range Plan for Nuclear Science, http://science.energy.gov/~media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf.
- [5] U. Heinz et al., arXiv:1501.06477 [nucl-th].
- [6] L. D. McLerran and R. Venugopalan, *Phys. Rev. D* 49, 2233 (1994) [hep-ph/9309289].
- [7] L. D. McLerran and R. Venugopalan, *Phys. Rev. D* 49, 3352 (1994) [hep-ph/9311205].
- [8] J. Berges, K. Boguslavski, S. Schlichting, and R. Venugopalan, *Phys. Rev. D* 89, no. 11, 114007 (2014) [arXiv:1311.3005 [hep-ph]].
- [9] N. Mueller, S. Schlichting, and S. Sharma, arXiv:1606.00342 [hep-ph].
- [10] T. Schafer and F. Wilczek, *Phys. Rev. Lett.* 82, 3956 (1999) [hep-ph/9811473].
- [11] R. Rapp, *Acta Phys. Polon. B* 42, 2823 (2011) [arXiv:1110.4345 [nucl-th]].
- [12] A. Mocsy, P. Petreczky, and M. Strickland, *Int. J. Mod. Phys. A* 28, 1340012 (2013) [arXiv:1302.2180 [hep-ph]].
- [13] T. Matsui and H. Satz, *Phys. Lett. B* 178, 416 (1986).
- [14] S. Pratt, E. Sangaline, P. Sorensen, and H. Wang, *Phys. Rev. Lett.* 114, 202301 (2015) [arXiv:1501.04042 [nucl-th]].

APPENDIX D: NUCLEAR PHYSICS CASE STUDIES

The following case studies were submitted by the authors listed below soon after the Nuclear Physics Exascale Requirements Review to guide development of nuclear physics requirements and the text of this report.

D.1 Case Studies Addressing Nuclear Astrophysics

- D-5 Core-Collapse Supernovae
Sean Couch, Michigan State University, Lansing, Michigan; and Bronson Messer and W. Raph Hix, Oak Ridge National Laboratory, Oak Ridge, Tennessee
- D-7 Neutron Star Mergers
Chris Fryer, Los Alamos National Laboratory, Los Alamos, New Mexico; and Dan Kasen, Lawrence Berkeley National Laboratory, Berkeley, California
- D-10 Neutrino Flavor Physics
Huaiyu Duan, University of New Mexico, Albuquerque, New Mexico; and George M. Fuller, University of California-San Diego, San Diego, California
- D-14 Astrophysical Thermonuclear Explosions
Michael Zingale and Alan Calder, Stony Brook University, Stony Brook, New York

D.2 Case Studies Addressing Data/Experiments

- D-18 High-Fidelity Long-Term Simulations of Beam-Beam Effects in Particle Colliders
Yves Roblin, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Balsa Terzic, Old Dominion University, Norfolk, Virginia
- D-20 Electron Cooling Simulation
He Zhang, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Yves Roblin, Thomas Jefferson National Accelerator Facility, Newport News, Virginia
- D-23 Computational and Data Storage Needs for the Neutron Electric Dipole Moment Experiment at the Spallation Neutron Source
L. J. Broussard, Oak Ridge National Laboratory, Oak Ridge, Tennessee; and B. Plaster, University of Kentucky, Lexington, Kentucky
- D-26 sPHENIX Data Acquisition and Analysis
Martin L. Porschke and Christopher H. Pinkenburg, Brookhaven National Laboratory, Upton, New York
- D-29 STAR, from data taking to analysis
Jérôme Lauret, Brookhaven National Laboratory, Upton, New York; R. Jefferson Porter (Ed.), Lawrence Berkeley National Laboratory, Berkeley, California
- D-35 Precision Extraction of Quark-Gluon-Plasma Properties
Steffen A. Bass, Duke University, Durham, North Carolina
- D-40 RIB/NSCL Active Target Time Projection Chamber (AT-TPC)
Wolfgang Mittig and Yassid Ayyad, National Superconducting Cyclotron Laboratory, Michigan State University, Lansing, Michigan

- D-44 Data Acquisition and Online Analysis at FRIB
FRIB Data Acquisition Working Group; Robert Varner, corresponding author
- D-49 High-Performance Computing for FRIB Design, Operation and Future Upgrades
Dali Georgobiani, Mauricio Portillo, Marc Hausmann, Mikhail Kostin (kostin@frib.msu.edu), Georg Bollen, and Frederique Pellemoine, Facility for Rare Isotope Beams, Michigan State University, Lansing, Michigan
- D-54 GRETA Streaming
M. Cromaz, Lawrence Berkeley National Laboratory, Berkeley, California
- D-58 Gamma-ray Tracking for the GRETINA/GRETA Spectrometer
M. Cromaz, Lawrence Berkeley National Laboratory, Berkeley, California
- D-62 FRIB Spectrometers
D. Bazin and R. Zegers, National Superconducting Cyclotron Laboratory, Michigan State University, Lansing, Michigan; and A. Couture and S. Mosby, Los Alamos National Laboratory, Los Alamos, New Mexico
- D-65 CUORE
Tommy O'Donnell, Virginia Polytechnic Institute and State University, Blacksburg, Virginia
- D-69 The ALICE Experiment at the LHC
R. Jefferson Porter, Lawrence Berkeley National Laboratory, Berkeley, California
- D-74 Hall-D/GlueX at Jefferson Lab
Mark Ito and David Lawrence, Hall D, Thomas Jefferson National Laboratory, Newport News, Virginia
- D-78 CLAS12 Physics Data Processing
H. Avakian, G. Gavalian, V. Gyurjyan, and V. Ziegler, Thomas Jefferson National Accelerator Facility, Newport News, Virginia

D.3 Case Studies Addressing Nuclear Structure and Reactions

- D-82 Nuclei from First Principles
G. Jansen, Oak Ridge National Laboratory, Oak Ridge, Tennessee; J. Vary, Iowa State University, Ames, Iowa; and S. Gandolfi, Los Alamos National Laboratory, Los Alamos, New Mexico; Hai Ah Nam, Los Alamos National Laboratory, Los Alamos, New Mexico; E. Ng, Lawrence Berkeley National Laboratory, Berkeley, California; and P. Maris, Iowa State University, Ames, Iowa
- D-88 Few-Nucleon Reactions
S. Quaglioni, Lawrence Livermore National Laboratory, Livermore, California; and R. Schiavilla, Old Dominion University and Thomas Jefferson National Accelerator Facility, Newport News, Virginia; Petr Navrátil, Tri-University Meson Facility, University of British Columbia, Vancouver, British Columbia
- D-92 Electroweak Phenomena
J. Carlson, Los Alamos National Laboratory, Los Alamos, New Mexico; J. Engel, University of North Carolina, Chapel Hill, North Carolina; and R. Schiavilla, Old Dominion University, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; S.C. Pieper, Argonne National Laboratory, Argonne, Illinois

- D-97 Quantified Heavy Nuclei
N. Schunck, Lawrence Livermore National Laboratory, Livermore, California;
W. Nazarewicz, Michigan State University, Lansing, Michigan; S. Wild,
Argonne National Laboratory, Argonne, Illinois
- D-101 Dense Nucleonic Matter
J. Carlson and S. Gandolfi, Los Alamos National Laboratory, Los Alamos, New Mexico

D.4 Case Studies Addressing Cold Quantum Chromodynamics

- D-105 Hadron Structure
K. Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility, Newport News, Virginia; W. Detmold, Massachusetts Institute of Technology, Cambridge, Massachusetts; R. Gupta, Los Alamos National Laboratory, Los Alamos, New Mexico; Huey-Wen Lin, Michigan State University, Lansing, Michigan; Meifeng Lin, Brookhaven National Laboratory, Upton, New York; S. Syritsyn, Stony Brook University, Stony Brook, New York; and B. Joo, D. Richards, and F. Winter, Thomas Jefferson National Accelerator Facility, Newport News, Virginia
- D-110 Hadron Spectroscopy
Robert Edwards, Balint Joo, and David Richards, Thomas Jefferson National Accelerator Facility, Newport News, Virginia
- D-114 Nuclear Physics from Lattice QCD
William Detmold, Massachusetts Institute of Technology, Cambridge, Massachusetts; Martin Savage, Institute for Nuclear Theory, Seattle, Washington; Balint Joo and Frank Winter, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Kostas Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility, Newport News, Virginia
- D-120 Fundamental Symmetries
R. Gupta, Los Alamos National Laboratory, Los Alamos, New Mexico; T. Kurth, National Energy Research Scientific Computing Center and Lawrence Berkeley National Laboratory, Berkeley, California; S. Syritsyn, Stony Brook University, Stony Brook, New York; Balint Joo, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; Kostas Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Martin Savage, Institute for Nuclear Theory, Seattle, Washington
- D-127 The Gluonic Structure of Matter
David Richards, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; Martin Savage, Institute for Nuclear Theory, Seattle, Washington; Robert Edwards and Balint Joo, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Kostas Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility, Newport News, Virginia
- D-132 Configuration Generation
Balint Joo and Robert Edwards, Thomas Jefferson National Accelerator Facility, Newport News, Virginia; and Thorsten Kurth, NERSC and Lawrence Berkeley National Laboratory, Berkeley, California; William Detmold, Massachusetts Institute of Technology, Cambridge, Massachusetts; and Kostas Orginos, College of William and Mary and Thomas Jefferson National Accelerator Facility, Newport News, Virginia

D.5 Case Studies Addressing Hot Quantum Chromodynamics

- D-138 Quark production and equilibration processes in non-equilibrium QCD
Sören Schlichting and Björn Schenke, Brookhaven National Laboratory,
Upton, New York

- D-142 Phases and Properties of Hot-dense Strongly Interacting Matter
Swagato Mukherjee and Patrick Steinbrecher, Brookhaven National Laboratory,
Upton, New York

- D-146 Spectral Functions, Transport Coefficients and Hadron Properties from Lattice QCD
Olaf Kaczmarek, Bielefeld University, Bielefeld, Germany

- D-149 Modeling Heavy-Ion Collisions
Scott Pratt, Michigan State University, Lansing, Michigan; and Björn Schenke,
Brookhaven National Laboratory, Upton, New York

D.1 Case Studies Addressing Nuclear Astrophysics

Core-Collapse Supernovae

Lead Authors: Sean Couch (Michigan State University); Bronson Messer (Oak Ridge National Laboratory), and W. Raph Hix (Oak Ridge National Laboratory)

1. Description of Research

1.1 Overview and Context:

Core-collapse supernovae (CCSNe) are among the most extreme laboratories for nuclear physics in the universe. Stellar core collapse and the violent explosions that follow give birth to neutron stars and black holes and, in the process, synthesize most of the elements heavier than helium throughout the universe. Despite the key role CCSNe play in many aspects of astrophysics and decades of research effort, *we still do not understand the details of the physical mechanisms that cause these explosions*. This leaves frustratingly large uncertainties on many key aspects of our theoretical understanding of the universe and also makes it difficult to constrain uncertain nuclear physics with data from CCSNe. The leading theory for the CCSN explosion mechanism is the delayed neutrino heating mechanism wherein neutrinos radiating away from the nascent neutron star deposit a sufficient fraction of their energy behind the stalled shock to revive it and drive an explosion. This is a highly nonlinear, multi-physics and multi-scale problem, and the only means of making progress has been through the application of large-scale computing.

1.2 Research Objectives for the Next Decade:

Full 3D simulations using multiple progenitor stars span a wide range of initial masses (i.e., at least from 9 to 40 solar masses). Realistic simulations of these scenarios require advanced multi-group neutrino radiation transport. These simulations should also have spatial resolution (via AMR) fine enough to resolve the development of the magneto-rotational instability, i.e., scales on the order of 100 meters within simulations spanning roughly 50,000 km in radius. To fully understand the production of heavy elements in such events, the target simulations should evolve nuclear reaction networks with as many as 3,000 isotopes in situ.

2. Computational and Data Strategies

2.1 Approach:

Simulations of the CCSN mechanism involve 3D general relativistic magnetohydrodynamics solvers coupled with multi-energy, multi-species neutrino transport, complex dense-matter equations of state, and detailed nuclear kinetic networks. Codes typically rely on domain decomposition for parallelism wherein different regions of the star are computed on different parallel devices (nodes, cores, etc.). Modern approaches add node-level parallelism to this mix, usually using OpenMP and GPU programming models to expose parallelism in local physics, i.e., neutrino transport, EOS, nuclear kinetics, etc. CCSN simulations generally stress I/O systems as well, requiring enormous amounts of data to be output at frequent intervals during the course of the simulations. In the next decade, codes tackling this problem will need to better exploit extreme levels of parallelism (i.e., accelerators), nonsynchronous communication patterns, and task-based parallelism.

2.2 Codes and Algorithms:

Many different codes are currently used to simulate CCSNe with high-fidelity (Chimera, FLASH, Zelmani, Prometheus-VERTEX, etc.). These codes are characterized by solving mixed sets of hyperbolic, parabolic, and elliptic PDEs using both explicit and implicit methods.

3. Current and Future HPC Needs

Core Collapse Supernovae Codes: Chimera/FLASH	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 2 (As a factor of column 1)
Computational core hours (conventional)	100M	0.1x	0.1x
Computational node hours (homogeneous many-core)	0M	10M	10x
Computational node hours (with GPU or accelerator)	10M	10x	100x
Memory per node	16 GB	16x	128x
Aggregate memory	60 TB	10x	100x
Data read and written per run	1,000,TB	20x	100x
Maximum I/O bandwidth needed	100 GB/sec	10x	40x
Percentage of runtime for I/O	1–2%	1–2%	1-2%
Scratch file system space needed	3,000 TB	20x	40x
Permanent online data storage	3,000 TB	20x	40x
Archival data storage needed	3,000 TB	20x	40x

Neutron Star Mergers

Lead Authors: Chris Fryer (Los Alamos National Laboratory) and Dan Kasen (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context:

The merger of two neutron stars is believed to produce short-duration gamma-ray bursts, to be a dominant source of rapid neutron-capture nuclear yields, and to be one of the main sources for gravitational waves. With the first detection of gravitational waves, the potential of using joint gravitational wave and electromagnetic wave emission to constrain properties of dense nuclear matter has moved from a theoretical possibility to a real scientific program. To actual use these mergers as probes (and understand their potential for gamma-ray bursts and nucleosynthesis), detailed models of the merger process and the electromagnetic emission are essential.

1.2 Research Objectives for the Next Decade:

Although models of neutron star mergers exist in the literature, these models have all made simplifying assumptions in the physics: either simplifying the effects of general relativity, the behavior of matter at nuclear densities (equation of state), or the neutrino transport (cross-sections or radiation transport). Magnetic fields are also important, and although a fully resolved, full-system, magnetohydrodynamic calculation is probably beyond what can be accomplished in the next decade, simulations should be able to advance our understanding of this physics considerably.

To tie these merger calculations to observations, we must model the electromagnetic radiation arising from both the relativistic jets and radioactive decay of the ejecta in these mergers. The computational physics problem requires implementing detailed line-dominated opacities into higher order transport codes, touching many of the current difficulties in radiation transport.

2. Computational and Data Strategies

2.1 Approach:

Both Eulerian (adaptive mesh refinement, AMR) and Lagrangian (smooth particle hydrodynamics, SPH) methods are employed with varying levels of physics. This physics will be improved over the next decade, presumably using asynchronous approaches.

2.2 Codes and Algorithms:

SPH calculations (e.g., SNSPH) use a hashed oct-tree method for the framework and asynchronous physics implementations. AMR methods vary, but most use block-based methods. Transport methods vary, but explicit methods parallelize better and will probably be best suited for advanced architectures.

3. Current and Future HPC Needs

3.1 Computational Hours:

For most of the merger codes (e.g., SNSPH, Cactus, “HAD”), typical current runs are 100,000 to one million core hours using conventional (MPI) methods with a factor of a few speed-up with homogeneous many-core runs. Post-process transport is relatively cheap with most runs at current resolutions running 100,000 core hours. If we need to couple these transport calculations to the hydrodynamics, the compute time will increase at least tenfold.

The usefulness of GPUs depends upon the communication time to these GPUs (improving) and whether there is a low-memory aspect of physics that can be put on these GPUs.

3.2 Parallelism:

Most codes use both MPI and Open-MP for parallelization and scale well on multi-processor and multi-core (relatively homogeneous) machines. Most of the advances are focused on additional physics (typically increasing floating point operations more than communication needs). Few codes are set to run on GPUs at present.

3.3 Memory:

The memory requirements for the merger simulations are proportional to the total resolution, and we expect to continue to increase the resolution with time, increasing the total memory (however, the memory per node may not increase dramatically). Fully general relativistic calculations require roughly a factor of 100 per cell/particle more memory.

3.4 Scratch Data and I/O:

Current computational constraints on these multi-physics simulations mean that we are not performing high-resolution runs with a full set of physics. This means that both the scratch storage and I/O bandwidth requirements are minimal, with ~1–2% runtime spent on I/O bandwidth. The total storage will probably remain below ~10 TB per run over the next 5–10 years.

3.5 Long-term and Shared Online Data:

The simulation results from merger calculations are used by many groups, and typically the data requirements do not decrease from the scratch data for long-term and shared data (approximately 1 TB per run increasing to 10 TB per run in 2025).

3.6 Archival Data Storage:

Currently roughly 100 TB of data are on archive from current runs. This will increase 100-fold by 2025.

3.8 Many-Core and/or GPU Readiness:

Most codes are ready for many-core simulations, but little has been done for GPU readiness.

Requirements Summary Worksheet

Neutron Star Mergers Code: CACTUS/HAD/SNSPH	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 2 (As a factor of column 1)
Computational core hours (conventional)	1M	10	50
Computational node hours (homogeneous many-core)	300K	10	50
Computational node hours (with GPU or accelerator)			
Memory per node	2 GB	2GB	2GB
Aggregate memory	0.2–2 TB	0.5–5 TB	1–10 TB
Data read and written per run	1 TB	5 TB	10 TB
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O	1–2%	2–4%	2–4%
Scratch file system space needed	0.1–1 TB	0.25–2.5 TB	0.5–5 TB
Permanent online data storage	0.1–1 TB	0.25–2.5 TB	0.5–5 TB
Archival data storage needed	0.1–1 TB	0.25–2.5 TB	0.5–5 TB

Neutrino Flavor Physics

Lead Authors: Huaiyu Duan (University of New Mexico) and George M. Fuller (University of California-San Diego)

1. Description of Research

1.1 Overview and Context:

For realistic studies of neutron star mergers and core-collapse supernovae, we study how neutrinos in different flavors and types (right- and left-handed) evolve in hot and dense astrophysical. In principle, neutrino flavor and/or spin transformation, sometimes loosely referred to as “neutrino oscillations,” must be followed with a full quantum kinetic treatment involving coherent neutrino propagation and scattering-induced flavor conversion. In some astrophysically important cases, however, it suffices to follow only the coherent component of flavor evolution. With or without scattering, this problem is inherently nonlinear because neutrinos and the flavors they carry determine the potentials that govern how neutrinos change their flavors. The computational crux of the neutrino flavor evolution problem is the necessity of following high-frequency quantum mechanical flavor phases. Given the inherent nonlinearity and the high-frequency-phase issues, numerical approaches are essential, although some (semi-) analytical solutions to simplistic models can provide useful physical insights. High-performance computing is necessary for solving full quantum kinetic/neutrino oscillation problems, and large storage space is needed for saving the data for later analysis.

1.2 Research Objectives for the Next Decade:

In the past decade, we have mostly focused on following coherent neutrino oscillations in a stationary, spherically symmetric supernova model, aka the bulb model. A series of recent studies has pointed out the insufficiency of this model. The scientific goals in the next decade are to determine how neutrino fluxes evolve in a realistic supernova environment and the impacts they may have on other aspects of supernova physics, including the explosion mechanism and nucleosynthesis. The computational goals are to develop and solve a multi-dimensional model, including various aspects of the physics, such as neutrino scattering and spin, that have been left out of most large-scale coherent bulb-like treatments.

2. Computational and Data Strategies

2.1 Approach:

The essence of the quantum kinetic neutrino-flavor transport problem in a hot and dense environment is to solve for the evolution of the (3x3 or 6x6) neutrino flavor density matrices. This problem comprises many coupled nonlinear 7D (time + space + momentum) Liouville-von Neumann equations. The current approach is to solve the coherent bulb model as an initial condition problem with given neutrino properties on the neutrino sphere. We expect a dramatic change to this approach in the next decade. For example, the neutrino fluxes must be solved with a boundary condition when the first-order scattering effects, i.e., “halo fluxes,” from within the supernova envelope are included. We also expect to include the temporal changes and/or more spatial dimensions into the problem.

2.2 Codes and Algorithms:

We have several numerical codes, most of which are based on the algorithms described in Duan et al. (2008). Because of the high degree of symmetry of the bulb model, neutrino fluxes can be described by thousands to tens of thousands of neutrino (angular) beams, each with hundreds to thousands of neutrino energy bins, and all of the neutrinos are assumed to be emitted from the same point on the neutrino sphere. Given the initial conditions on the neutrino sphere and the density profile, the neutrino beams are evolved simultaneously over radius by solving millions to tens of millions of coupled nonlinear differential equations with a modified midpoint method, and snapshots are saved along the way. We also employ a numerical code, BURST, which follows Boltzmann neutrino transport coupled to nucleosynthesis in the early universe.

3. Current and Future HPC Needs

3.1 Computational Hours:

The computational hours and resources of current HPC needs listed in this section are based on 100 runs of XFLAT, each of which corresponds to one of the many snapshots during a full supernova simulation. XFLAT is a C++ code recently developed at UNM to explore the Intel Many Integrated Core (MIC) architecture. It can run on both the homogenous HPC platform with the CPU or the Xeon Phi accelerator only and the hybrid platform with both the CPU and Xeon Phi. However, this code does not include some of the physics, such as the three-flavor mixing that has already been implemented in other codes.

It is difficult to estimate the resources that may be required in the next decade for solving the problem of neutrino oscillations near compact objects. The estimates of the HPC needs in 2020 and 2025 are based on one potential research direction, i.e., solving neutrino oscillations in a multi-D supernova model as an initial condition problem (without neutrino halo fluxes) by 2020, and solving the multi-D model with neutrino halo fluxes by 2025. These estimates do not include, e.g., the study of the impacts of neutrino oscillations on nucleosynthesis. A key objective will be to upgrade BURST to a full quantum kinetic treatment of neutrino flavor physics coupled to nucleosynthesis in the early universe. This will guide us in a similar treatment in compact objects.

3.2 Parallelism:

Following neutrino-flavor evolution in compact objects and the early universe lends itself nicely to parallel treatment. XFLAT uses all three levels of parallelism in the current HPC paradigm. It distributes neutrino beams among the computing nodes and computing devices (i.e., CPU and/or Xeon Phi) within each computing node. The computing devices are synchronized through MPI. Within each computing device, neutrino beams are further distributed among cores via OpenMP threads (~16 threads on a CPU and ~240 threads on a Xeon Phi). On each core, the operations on the neutrino energy bins within each beam are carried out by utilizing SIMD (with the width of 4–8 double precision numbers). Typically, fewer than 50 nodes are requested in each run given the modest computing needs of the Bulb model. However, it is almost certain that this number will increase significantly in the next decade after a multi-D model is developed.

3.3 Memory:

The memory requirement of the problem is modest compared to data intensive HPC problems. We have no plan to micromanage the memory.

3.4 Scratch Data and I/O:

XFLAT does not demand much scratch space and I/O bandwidth at the moment, but much more scratch space may be needed in the next decade. This consideration also applies to BURST and the quantum kinetic code that may follow.

3.5 Long-term and Shared Online Data:

XFLAT does not demand much online storage at the moment, but much more may be needed in the next decade.

3.6 Archival Data Storage:

XFLAT does not demand much archival storage at the moment, but much more may be needed in the next decade.

3.7 Workflows:

In a current typical numerical run, we choose certain time snapshots from a full supernova simulation. Using the neutrino emission parameters (energy spectra, luminosities, radii of the neutrino sphere, etc.) and matter profile of each snapshot, we compute and save the neutrino fluxes in all flavors throughout the supernova. The stored results are later used to study the impacts of neutrinos on supernova nucleosynthesis and neutrino signals. Likewise, in BURST and in planned quantum kinetic follow-ons, we calculate neutrino occupation probabilities for binned energy distributions for all flavors and nuclear abundances at each time step in the early universe.

3.8 Many-Core and/or GPU Readiness:

As explained in 3.2, XFLAT is ready for HPC with many CPU-like cores, heavy or light. BURST is likewise being configured to take advantage of this.

3.9 Software Applications, Libraries, and Tools:

We use standard software packages, such as the Intel C++ compiler with MPI, OpenMP, and SIMD support, as well as HDF5 and NetCDF.

Reference

H. Duan, G. M. Fuller, and J. Carlson. *Comput. Sci. Disc.* 1 (2008) 015007.

Requirements Summary Worksheet

Coherent Neutrino Oscillations Code: XFLAT	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 2 (As a factor of column 1)
Computational core hours (conventional)	40,000	x100	x1,000
Computational node hours (homogeneous many-core)**	1,200	x100	x1,000
Computational node hours (with GPU or accelerator)***			
Memory per node	<1 GB	64 GB	256 GB
Aggregate memory	<0.1 TB	1 TB	1,000 TB
Data read and written per run	1 TB	1,000 TB	10,000 TB
Maximum I/O bandwidth needed	1 GB/sec	10 GB/sec	1000 GB/sec
Percentage of runtime for I/O	5%	5%	5%
Scratch file system space needed	1 TB	1,000 TB	10,000 TB
Permanent online data storage	1 TB	1,000 TB	1,000 TB
Archival data storage needed	1 TB	1,000 TB	1,000 TB

Astrophysical Thermonuclear Explosions

Lead Authors: Michael Zingale and Alan Calder (Stony Brook University)

1. Description of Research

1.1 Overview and Context:

Many explosive astrophysical phenomena are powered by runaway thermonuclear reactions occurring in material in the interior or on the surface of stars. These events include classical novae, type I X-ray bursts (XRBs), and type Ia supernovae (SNe Ia)—all involving matter transferred from a companion onto a compact object giving rise to a thermonuclear runaway. These events produce and disseminate heavy elements. Modeling these events requires a multi-physics, multi-scale hydrodynamics code, where the additional physics, gravity and nuclear reaction networks, and the scales of these problems sometimes necessitate subgrid-scale models for unresolvable physics like flames and background turbulence. These simulations require the use of the largest extant computing platforms, and even on these, we cannot approach the range of length scales at play in the real stellar system. Algorithm development is a critical part of the process of modeling astrophysical explosions.

1.2 Research Objectives for the Next Decade:

There are several objectives for the next decade: (1) Continuing the process of understanding the science of thermonuclear explosions. At present, we are investigating the role of convection in type I X-ray bursts, the Urca process in the progenitors of SNe Ia, convection and ignition in the sub-Chandra model of SNe Ia, and the efficacy of merging white dwarfs as the SNe Ia progenitors. These projects will continue for the next few years. (2) Pushing the size of the problems we study, including increasing the spatial resolution, to better capture the burning and turbulence at small scales, and larger domains (for the XRBs). (3) Improving the physics, in particular, larger reaction networks to capture the nucleosynthesis, and adding magnetic fields and rotation. (4) Improving the performance of our application codes. Our present effort is centered on software development to port the modules for nuclear burning and the equation of state to Graphics Processing Units (GPUs). We will also explore better load balancing and task management, in situ diagnostics and visualization, and asynchronous physics integration.

2. Computational and Data Strategies

2.1 Approach:

In practice, we can never attain the resolution required to capture all of the physical scales at play in our systems, so we model the largest scales that are important and use a subgrid model (or do ILES) to represent the physics at unresolved scales. This means that our goal with larger systems will always be pushing to bigger systems, to increase the realism of our simulations. Our data output is large—often 100 TB per simulation (excluding checkpoint files). We do our analysis on the remote machines today, but with larger machines, we will need to switch to in situ analysis.

2.2 Codes and Algorithms:

The two codes in use are Castro and Maestro. Castro is a fully compressible radiation hydrodynamics code, while Maestro is a low Mach number hydrodynamics code—able to efficiently model reacting, convective flows in stars. The two codes complement one another, allowing for simulations to begin in the slow simmering convective phases that often precede thermonuclear runaways and follow the evolution into the explosive phase. All of our codes are freely available on github (<https://github.com/BoxLib-Codes>). The two codes are built on top of the BoxLib adaptive mesh refinement library.

Maestro solves the equations of hydrodynamics under the assumption of a low Mach number. This is expressed by requiring that the pressure everywhere in our domain remain close to the background hydrostatic pressure of the star we are modeling. Mathematically, this constraint manifests itself as an elliptic constraint on the velocity field; physically, the effect is to filter sound waves from the system. Together with a hyperbolic system of PDEs describing conservation of mass, momentum, and energy, we can efficiently model convective flows in a stratified background. We use a finite-volume discretization to solve the equations, advancing the state in each time step using a fraction step method consisting of reactions and advection, and projecting the velocity field to satisfy the constraint. The projection takes the form of a variable-coefficient Poisson problem. Maestro allows for time steps on the fluid velocity time scale, rather than the sound speed, allowing for much longer time integration.

Castro solves the fully compressible hydrodynamics equations coupled to a multi-group flux-limited diffusion implementation of radiation. It uses an unsplit Godunov method for the advective terms and includes self-gravity and rotation in a conservative flux framework. Both codes share the same microphysics (stellar equation of state and reaction networks). BoxLib manages the grid patches and parallelization and uses a hybrid MPI + OpenMP approach to make efficient use of modern architectures, and the enforcement of Maestro's constraint is done using the BoxLib geometric multigrid solvers.

3. Current and Future HPC Needs

3.1 Computational Hours:

At present, our computation takes place entirely on the CPUs, although we have an active effort porting our microphysics (equations of state and reaction networks) to the GPUs. We currently use about 50 M core-hours per year for our science simulation, roughly split between the two application codes.

(Note: The numbers in the Requirements table below are for Maestro. The numbers for Castro applications will be similar.)

3.2 Parallelism:

Castro and Maestro use a hierarchical parallelization strategy to run on massively parallel supercomputers. These achieve coarse-grained parallelism by the distribution of rectangular patches of data to processors with communication between processors using the Message Passing Interface (MPI) library. Additional fine-grained parallelism is achieved using the OpenMP library, which allows multiple processors to work on the same patch of data at the same time. Both codes have demonstrated excellent strong scaling out to approximately 10,000s of cores and weak scaling to 200,000 cores. Threading in Castro is now implemented using tiling and we will extend this approach to Maestro in the near future.

3.3 Memory:

We typically don't run into memory problems with our simulations—with the hybrid approach to parallelism, we can make efficient use of the shared memory on a processor. We expect to continue this trend in the future.

3.4 Scratch Data and I/O:

The more disk space, the better, but we can do our work flows with 10–20 TB of project disk space. We can then script the retrieval of data from HPSS and process it in chunks. Our I/O takes a small fraction of our runtime (<1%) and performs efficiently (e.g., >10 GB/sec on Titan). We do not expect I/O to take more than 5% of our runtime in the next 10 years. Our codes have 2 levels of plot files to allow for smaller datasets to be written out at greater frequency.

3.5 Long-term and Shared Online Data:

We need approximately 100 TB *per simulation* of mass storage. We typically need to keep the data around for 1–2 years to see the analysis and publication process through. We have not needed to share data online yet, but can see the need in the future.

3.6 Archival Data Storage:

We have several 100s of TBs of data stored in archival storage (HPSS). We expect this number to double every few years.

3.7 Workflows:

Our codes are maintained via git and hosted on github. We do regular regression testing on the codes (nightly). Our simulations keep track of the git hashes for the codes used and store these in all output products. We do our analysis on the remote machines mainly using yt and python. We expect a similar work flow to carry forward over the next decade, but also see the need for runtime visualization.

3.8 Many-Core and/or GPU Readiness:

We have invested considerable development time to allow the codes to use GPUs on platforms such as Titan at OLCF. This will continue in the future. BoxLib uses tiling to efficiently operate on many-core architectures, and our codes benefit from the core optimizations done by the BoxLib group.

3.9 Software Applications, Libraries, and Tools:

We use standard libraries and programming models: MPI, OpenMP, OpenACC, HDF5, hypre.

Requirements Summary Worksheet

Astrophysical Thermonuclear Explosions Code: Maestro	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 2 (As a factor of column 1)
Computational core hours (conventional)	10,000	2.5	5
Computational node hours (Homogeneous many-core)	5M–10M h per simulation	5	5
Computational node hours (with GPU or accelerator)	in progress	5M–10M h per simulation	5
Memory per node	2–4 GB	5x GB	10x GB
Aggregate memory	500 TB	5x TB	10x TB
Data read and written per run	100 TB	5x TB	10x TB
Maximum I/O bandwidth needed	10 GB/sec	20 GB/sec	50 GB/sec
Percentage of runtime for I/O	1%	1–2x	2–4x
Scratch file system space needed	10 TB	5x TB	10x TB
Permanent online data storage	10 TB	5x TB	10x TB
Archival data storage needed	500 TB	1,000 TB	2,000 TB

D.2 Case Studies Addressing Data/Experiments

High-Fidelity Long-Term Simulations of Beam-Beam Effects in Particle Colliders

Lead Author(s): Yves Roblin (Thomas Jefferson National Accelerator Facility) and Balsa Terzic (Old Dominion University)

1. Description of Research

1.1 Overview and Context:

The simulation of beam-beam effects in particle colliders is a challenging task. Possible future machines, such as Jefferson Laboratory's electron-ion collider (JLEIC) are limited in luminosity by such effects. The interaction between two colliding beams is described by a Poisson equation, and a number of methods have been developed to solve it efficiently. However, the next generation of colliders is considering using a different number of bunches in each ring resulting in each bunch of one ring colliding with every other bunch of the other ring. This has several benefits ranging from efficient timing synchronization to better systematics in some physics measurements, such as the determination of beam polarization. This also drastically increases the computational load to the point where conventional methods become prohibitive.

1.2 Research Objectives for the Next Decade:

Our goals are to push the capabilities of beam-beam simulations codes to address ring-ring synchronization problems where gear changing occurs as well as open the door to long-term simulations (billions of turns, several hours of machine time) of the interactions to study the dependency of machine lifetime on a range of higher level parameters (beam polarization, magnet imperfections, and other quantities).

2. Computational and Data Strategies

2.1 Approach:

Existing codes such as **BeamBeam3D** are utilizing a shifted Green function approach to solving the Poisson equation, which, while accurate, is computationally expensive. We are developing a new code, **GHOST** to specifically address previously intractable long-term high-fidelity simulations by utilizing the computational prowess of the emerging hybrid CPU/GPU computing architectures. The new code features particle tracking symplectic to an arbitrary order, substantially more efficient than the Bassetti-Erskine approximation to a Poisson equation, optimized for performance on a cluster of GPU devices.

2.2 Codes and Algorithms:

GHOST utilizes a layered approach where, at the core of it, GPU units are used for the calculation of the interactions. These are then federated at the next layer by a high-speed networking transport and efficient partitioning of data flows. There is an important advantage of implementing our code on a parallel GPU platform over a parallel CPU platform. The increase in computational power of CPUs as a function of time has followed Moore's Law, before recently saturating due to physical limitations of cramming components on integrated circuits. A power law similar to Moore's Law has governed the

increase in computational power and memory bandwidth of GPUs, and is expected to continue in the foreseeable future. Therefore, the performance of codes particularly optimized to run on GPUs will continue to improve as they are executed on new generations of GPU architectures.

Requirements Summary Worksheet

Code: GHOST	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)			
Computational node hours (homogeneous many-core)			
Computational node hours (with GPU or accelerator)*	1,000	x10	X 100
Memory per node	16 GB	X8	X64
Aggregate memory	16 TB	X8	X64
Data read and written per run	1 TB	X8	X64
Maximum I/O bandwidth needed	288 GB/sec	X8	X64
Percentage of runtime for I/O			
Scratch file system space needed	TB	TB	TB
Permanent online data storage	- TB	10 TB	50 TB
Archival data storage needed	- TB	40 TB	200 TB

* Node hours for GPU or accelerator usage.

Electron Cooling Simulation

Lead Author(s): He Zhang and Yves Roblin (Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

Electron cooling is a technique to reduce the emittance and increase the intensity of ion beams. For future colliders, such as the JLab electron ion collider (JLEIC), electron cooling is the essential technique to compensate the intrabeam scattering (IBS) effect, to achieve and maintain the low emittance of the ion beam for the proposed high luminosity. The design of the electron cooler, and even the design and the optimization of the whole collider, heavily depend on numerical simulations.

1.2 Research Objectives for the Next Decade:

In future collider design, a high-energy (above 50 MeV) and high-density (up to 2 nC/bunch) bunched electron beam is proposed to be implemented in electron cooling. High magnetic fields (around 2T) are needed to mitigate the collective effect, e.g., space charge effect, of the electron beam. A thorough understanding of the physics in the cooling process at this parameter range is necessary and imperative. Simulations of the motion of the ions and electrons inside the cooler are required to understand the dynamic process of electron cooling at the microscopic level. Calculation of the interaction between billions or more charged particles requires significant HPC resources.

2. Computational and Data Strategies

2.1 Approach:

In the microscopic simulation on the electron cooling process, one first needs to accurately calculate the interaction between the ions and the electrons, and then simulate their motion by solving the dynamic equations for each particle. Collisions between the electrons and the ions have to be correctly modeled. The state-of-the-art code can only simulate the motion of a few ions inside an electron beam. In the future, one may want to simulate a real ion beam, or at least a better sample of it. One may also want to include the space charge effect and/or other collective effects in the simulations.

2.2 Codes and Algorithms:

To correctly model the collision, real particles, instead of macroparticles, are used in the simulation, which makes the field calculation very computationally expensive. Currently we are testing fast field solvers, e.g., the fast multipole method. Once the field is known, the dynamic equations can be solved to promote the particles. Nearregion collisions require a very small step size to model. Analytic formulas describing the two-body collisions are used to allow a moderate step size and an increase in simulation efficiency.

3. Current and Future HPC Needs

3.1 Computational Hours:

3.2 Parallelism:

Coarse-grained parallelization can be easily implemented in solving the dynamic equations by distributing the particles to the nodes. Parallelization of the Poisson field solver is more complicated. But some fast solvers, for example, the fast multiple method we are testing now, can take advantage on both the coarse-grained and the fine-grained parallelism.

3.3 Memory:

3.4 Scratch Data and I/O:

3.5 Long-term and Shared Online Data:

3.6 Archival Data Storage:

3.7 Workflows:

3.8 Many-Core and/or GPU Readiness:

Some fast multi-pole method package can use the accelerators (e.g., GPUs) in computation. Usage of the new hardware for the other part of the simulation needs more investigation.

3.9 Software Applications, Libraries, and Tools:

3.10 HPC Services:

3.11 Additional Needs:

Requirements Summary Worksheet

Code:	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	500,000	2X	5X
Computational node hours (homogeneous many-core)	500,000	2X	5X
Computational node hours (with GPU or accelerator)			
Memory per node	2 GB	2X col. 1 GB	10X GB
Aggregate memory	0.2 TB	2X col. 1 TB	10X TB
Data read and written per run	0.001 TB	2X TB	10X TB
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O			
Scratch file system space needed	0.5 TB	3X	6X
Permanent online data storage	1 TB	3X	6X
Archival data storage needed	2 TB	3X	6X

* Core hours for conventional processors.

** Node hours for homogenous many-core architectures.

Computational and Data Storage Needs for the Neutron Electric Dipole Moment Experiment at the Spallation Neutron Source

L.J. Broussard (Oak Ridge National Laboratory) and B. Plaster (University of Kentucky)

1. Description of Research

1.1 Overview and Context:

The discovery of a neutron electric dipole moment (EDM) would provide unambiguous evidence for new beyond Standard Model T- and CP-violating physics. A new search for the neutron EDM is under way at the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory. The goal of this SNS nEDM experiment is to achieve a sensitivity of $<5 \times 10^{-28}$ e-cm, which is two orders of magnitude below the present limit. Large components of the apparatus are projected to begin commissioning on the floor of the SNS beginning in 2020, and final assembly and commissioning and the beginning of data collection are projected to begin in 2022. Control of systematic errors is of paramount importance to nEDM experiments; several of these can be studied via detailed calculations and simulations.

1.2 Research Objectives for the Next Decade:

During the next decade, the worldwide neutron EDM community is poised to achieve at least a one order-of-magnitude improvement in sensitivity beyond the present limit of $<3 \times 10^{-26}$ e-cm. During this same timeline, the SNS nEDM experiment will have begun data collection for its projected two orders-of-magnitude improvement in sensitivity. Ultimately achieving the final result will require detailed study of systematic errors. Several of these can be studied using “traditional” simulation packages, such as a Geant4; however, other codes, such as COMSOL, will be employed for their electromagnetic and spin-tracking capabilities. Ultimately, the experiment’s goal is to have carried out detailed simulations of all (known) systematic effects to the level of the experiment’s sensitivity, and to have in place capabilities for real-time, online waveform analysis of the data stream (for example, using GPUs).

2. Computational and Data Strategies

2.1 Approach:

Simulations of many of the systematic errors in the nEDM experiment tend to be highly computationally intensive. Several examples include detailed models for neutron spin tracking/integration in electromagnetic field gradients (notoriously, a computational challenge), simulations of beam-generated backgrounds (requiring very high statistics), and other quantities. At present, collaborators have mostly executed their simulation codes on local, small-scale clusters. The experiment would benefit significantly from access to a centralized cluster with large-scale storage (access to these local clusters is typically restricted to local users only), permitting the generation of large-scale simulated data sets and facilitating collaborative access.

2.2 Codes and Algorithms:

To date, the collaboration has employed Geant4 for simulations of neutron beam transport, neutron activation backgrounds, neutron spin tracking, and “detector response” (including optical transport simulations) to various event types, including backgrounds. COMSOL has been employed for detailed finite-element analysis calculations of electromagnetic fields, and also for spin-tracking, in an alternative to Geant4.

3. Current and Future HPC Needs**3.1 Computational Hours:**

The computational hours currently employed on conventional cores for our Geant4 and COMSOL calculations are summarized in their respective tables below. As the experiment moves towards commissioning and data-taking, the projected required computational hours will increase significantly.

3.2 Parallelism:

At present, only the COMSOL codes are parallelized. The Geant4 code is “parallelized” in the sense that separate jobs are stitched together to combine statistics. We will certainly move towards improving this in the future, especially with regard to large-scale Geant4 simulations.

3.3 Memory:

The memory requirements for the COMSOL code (which can be quite substantial) are summarized in its table below. Current efforts are limited by insufficient total memory.

3.4 Scratch Data and I/O:

We anticipate requiring up to approximately 100 TB of scratch data space by 2020–2025 for simulations and data analysis. We do not anticipate any extraordinary I/O needs.

3.5 Long-term and Shared Online Data:

We also anticipate requiring up to approximately 1 PB of long-term and shared online data during the intensive data analysis phase of the experiment.

3.6 Archival Data Storage:

We anticipate requiring up to approximately 1 PB of archival data storage (e.g., waveforms).

3.7 Workflows:**3.8 Many-Core and/or GPU Readiness:**

At present, our collaboration’s codes are not at the frontier of many-core and/or GPU readiness. Of course, we would like to reach this level. This will require the investment of human resources, such as additional funding for postdoctoral scholars with significant background in this area, to work on simulations and data analysis.

3.9 Software Applications, Libraries, and Tools:

We do not anticipate any unusual needs here.

3.10 HPC Services:

We do not anticipate requiring any extraordinary HPC services in 2020–2025, other than standard support for data storage, transfer, etc.

3.11 Additional Needs:**Requirements Summary Worksheet**

Code: Geant4	Column 1: CurrentUsage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	$10^5/\text{year}$	$\times 10^3$	$\times 10^4$
Computational node hours (homogeneous many-core)			
Computational node hours (with GPU or accelerator)			
Memory per node	32 GB	32 GB	32 GB
Aggregate memory	0.125 TB	30 TB	300 TB
Data read and written per run	0.01 TB	0.01 TB	0.01 TB
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O	~ 0	~ 0	~ 0
Scratch file system space needed	0.5 TB	10 TB	100 TB
Permanent online data storage	10 TB	100 TB	1000 TB
Archival data storage needed	0 TB	100 TB	1000 TB

sPHENIX Data Acquisition and Analysis

Lead Author(s): Martin L. Purschke and Christopher H. Pinkenburg (Brookhaven National Laboratory)

1. Description of Research

1.1 Overview and Context:

The existing PHENIX experiment will end its data-taking in the summer of 2016 and will be replaced with an ambitious upgrade, code-named “sPHENIX.” The new experiment will be centered around the former BaBar superconducting solenoid and consist of an inner tracking system, an electromagnetic calorimeter, and 2 layers of hadronic calorimeters. The upgrade will give the experiment full azimuthal coverage within a pseudorapidity range of $-1.1 < \eta < 1.1$. The new experiment is expected to be installed and commissioned in 2021 and begin data-taking in 2022.

1.2 Research Objectives for the Next Decade:

The new experiment is designed to provide a comprehensive measurement of jets and Υ mesons, which are the next frontiers in heavy-ion physics and also lighter collision systems. The computing resources required are needed for simulations to find the optimal detector design, the acquisition and long-term storage of the data, and the reconstruction and analysis of the data. The data volume is estimated at 30 PB/year. From the experience from the PHENIX experiment, $\frac{1}{4}$ of the raw data size will be needed for the Data Summary Tapes (DSTs), which are the main input the various analysis projects.

2. Computational and Data Strategies

2.1 Approach:

The computing efforts can be divided into simulations, storage, and data analysis. Current ongoing analysis of the PHENIX data runs in the framework of so-called “analysis taxi,” in which user jobs accessing the same data sets are bundled and re-use the once-retrieved data with high efficiency. All data files are kept on tape in the HPSS system for long-term storage. The most frequently accessed files, typically the DSTs, are stored in dCache. We do not expect a technology to emerge that would be able to outperform this computing paradigm, which has become routine and has very streamlined setup procedures.

2.2 Codes and Algorithms:

The sPHENIX simulations all use Geant4. The code to describe the detector components is managed in github. Both the simulations and the data analysis are done in root, which provides standard tools, such as histograms and DST storage, and also adds a convenient macro input layer to the simulations.

3. Current and Future HPC Needs

3.1 Computational Hours:

Based on about 13,000 batch slots, we estimate about 60 million CPU hours per year. Some additional peak demands will emerge. In the face of ever-more-sophisticated simulations requirements and the increase in acquired data volumes, this number is expected to go up by a factor of 8.

3.2 Parallelism:

At this time, we have virtually no applications making use of parallel architectures, such as GPUs.

3.3 Memory:

The current standard value of 2 GB/core has become a limit especially for simulation jobs. We expect the number of cores and the per-core memory to at least double.

3.4 Scratch Data and I/O:

Much of our scratch space is used by jobs on the local node and can remain local. We currently use about 400 TB of what we term “workspace,” which is not job-related, e.g. to collect the output from many jobs.

3.5 Long-term and Shared Online Data:

We are recording about 1 PB/year. We are keeping the most-wanted DSTs online in dCache. Most of the raw data are accessed locally (e.g., no GRID). This is expected to change. Reconstruction jobs accessing the raw data will always have priority locally to avoid wide-area data transfers.

3.6 Archival Data Storage:

The envisioned data rate for 2022 (when the experiment is expected to start taking data) is 30 PB/year, although this number has some uncertainties and depends on the eventual detector configuration, which is still in flux.

3.7 Workflows:

Unless a much higher-performing technology emerges, we will continue to use the current technology outlined in 2.1. Distributed computing resources (such as, but not confined to, the one provided by the Open Science Grid) will start to play a much larger role.

3.8 Many-Core and/or GPU Readiness:

There has been some progress using GPUs for simulations, but we expect the bulk of simulations and all of the data analysis to take place on CPUs.

3.9 Software Applications, Libraries, and Tools:

We have added our own libraries to the otherwise standard tools: root, Geant4, dCache, and others.

3.10 HPC Services:

The most imminent need of an HPC service would be the presence of convenient and manageable single-sign-on system to authenticate users.

3.11 Additional Needs:

Requirements Summary Worksheet

Code: Analysis/Reconstruction	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	65M	8X	12X
Computational node hours (homogeneous many-core)**	10M	6X	10X
Computational node hours (with GPU or accelerator)	--		
Memory per node	16 GB	64 GB	128 GB
Aggregate memory	30 TB	120 TB	240 TB
Data read and written per run	1 PB	30 PB	50 PB
Maximum I/O bandwidth needed	12 GB/sec	40 GB/sec	60 GB/sec
Percentage of runtime for I/O			
Scratch file system space needed	400 TB	800 TB	1200 TB
Permanent online data storage	2 PB	16 PB	24 TB
Archival data storage needed	10 PB	100 PB	150 PB

* Core hours for conventional processors.

** Node hours for homogenous many-core architectures.

STAR, from Data Taking to Analysis

Lead Author: Jérôme Lauret (Brookhaven National Laboratory)

Edited by: R. Jefferson Porter (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context:

The STAR experiment has phased in several major detector upgrades to enhance its detection capabilities in the mid-rapidity region and is well suited for a broad range of Heavy Ion and Spin studies at the Relativistic Heavy-Ion Collider (RHIC). STAR's two recent upgrades, the Heavy Flavor Tracker (HFT) for open heavy-flavor measurements and the Muon Telescope Detector (MTD) for quarkonia measurements, in operation for a few years, are bringing new unprecedented insights to the Heavy Flavor program. Future upgrades include the earmarked iTPC upgrade (Inner TPC), an Event Plane, and the later EndCap-Time-of-Flight (eTOF) detector proposed by an international STAR/CBM-TOF collaboration (Heidelberg, Darmstadt, Tsinghua, CCNU, and USTC). These changes will enable the collection of data that is critical to the physics mission for Phase II of the Beam Energy Scan (BES II).

With its massive amount of data taken each year (10s of PBytes/year), STAR has developed many strategies to deal with a demanding and versatile program. Some of the processing remains located at our Tier-0 center (RCF) and close to the experimental apparatus, but a significant fraction of the data production and all simulations are handled on distributed resources, such as that aggregated by the Open Science Grid (OSG).¹ Although the resources accessed in this manner are pledged to STAR, the software stack provided by the OSG is a critical component of our processing pipeline. Data sets are routinely transferred to our Tier-1 facility (NERSC), and in the coming years, increasing resource demands will likely cause additional traffic to Tier-2 analysis centers. This trend has already begun as storage space is insufficient at this time to accommodate all possible research topics.

1.2 Research Objectives for the Next Decade:

With its upgrades, the STAR collaboration will be in a position to tackle many challenging key scientific questions and harvest unprecedented scientific discoveries. A few of the possible objectives are highlighted herein. Through measurements of single spin asymmetries in $W^{+/-}$, Z , direct photon, and Drell-Yan production in transversely polarized $\sqrt{s} = 500$ GeV p+p collisions, the STAR detector would make the first significant measurement of the sign change of the Sivers function, a measurement representing a fundamental test of QCD. STAR also aims at clarify the interpretation of measurements related to the chiral magnetic effect, chiral magnetic wave, and chiral vortical effect and in this respect has proposed an original Beam Use Request colliding isobaric nuclei (^{96}Ru and ^{96}Zr). In 2019–2020 the STAR Collaboration proposes a second phase of the Beam Energy Scan program at RHIC (BES Phase-II) to answer compelling questions about the phase structure of QCD matter that cannot be addressed using existing measurements. The BES Phase-II program, with the upgrades discussed above, will allow for high-statistics measurements, with an extended kinematic range in rapidity and transverse momentum, using sensitive observables, to reveal the structure of the QCD phase diagram.

The 2015 NSAC Long Range plan notes in the first recommendation that, *“the upgraded RHIC facility provides unique capabilities that must be utilized to explore the properties and phases of quark and*

¹ The Open Science Grid consortium, URI=<https://www.opensciencegrid.org/>

gluon matter in the high temperatures of the early universe and to explore the spin structure of the proton.” Beyond 2020, STAR proposes several key measurements to complete its RHIC mission that are complementary to sPHENIX at RHIC as well as experiments at the LHC and JLAB. Those key physics opportunities include

- Studies of the nuclear parton distribution and fragmentation functions
- Understanding of the nature of the pomeron and potentially discovering the odderon
- Extension of gluon polarization results down to low x , to make a measurement of the gluon helicity distribution $\Delta g(x)$
- The evolution of transverse-momentum dependent distribution functions
- Constraints on the transport coefficients near the critical temperature T_c .
- Constraints on the 3+1D hydrodynamic and temperature dependence of QGP properties
- Open beauty and jet measurements

Many of the above measurements will be possible with the current and outlined upgrades and the capabilities of the current detector setup. But the full suite of measurements will require cost effective upgrades, such as a modest forward tracking and calorimeter upgrade, as well as an enhanced HFT tracking detector. This program would dramatically help the community advances toward the Electron-Ion-Collider (EIC), the future of the Nuclear Physics program.

2. Computational and Data Strategies

2.1 Approach:

The current data work flow includes Data Acquisition with the writing of the data stream onto “event buffer” and Meta-Data collection. The data is pushed to a Mass Storage system (HPSS) where it is available for later Offline processing, carried through a campaign that usually begins about 2 months after the data-taking has ended. The produced data is (a) copied back to Mass Storage, (b) moved to the Scalla/Xrootd² distributed storage space, and (c) copied to our Tier-1 NERSC facility. Online collection of conditions is done at a rate varying from 150 to 2,000 messages/sec by a Message Queuing collector known as MIRA,³ with 1,900 meta-data variables being continuously monitored and displayed in real time. In total, we have a little over three billion messages passing through the system per RHIC run.

Data analysis has been based on an “anyone can do any analysis anytime” model, which has been one of STAR’s main strengths. With recent budget limitations for growing the available storage, STAR is considering a strategy around a “Dataset Carousel” concept, based on maximizing storage availability by rotating data sets. Further out in time, the derived data (Data Summary Tape or DST), currently in form of Micro-DST (or MuDST) will likely evolve toward more compact format that could reduce the data output demand by a factor of 5 to few PB/year, transferable over the WAN. By the next decade, the transfer of smaller (picoDST and Trees) data sets would become routine from BNL or our Tier-1 center(s) to our Tier-2 analysis centers. At the Tier-0 center (BNL) the analysis will continue to make use of distributed storage (8 PB total) served by Scalla/Xrootd. The Dataset Carousel⁴ strategy would also bring

² XRootD URI=<http://xrootd.org/>

³ D. Arkhipkin and J. Lauret, “STAR Online Framework: from Metadata Collection to Event Analysis and System Control”, ACAT 2014 proceedings, *J. Phys.: Conf. Ser.* **608** 012036 [doi:10.1088/1742-6596/608/1/012036](https://doi.org/10.1088/1742-6596/608/1/012036) (2015)

⁴ Lauret, J. and Yu, D. “ERADAT and Data Carousel systems at BNL: A tool and UI for efficient access to data on tape with fair share policies capabilities”, *ACAT 2010*, *PoS ACAT2010 023* (2010)

an immediate relief to the currently unsustainable situation and may very well restore, with modest commitment, a breadth of physics opportunities and scientific endeavors. While storage is at this time the direst issue for STAR, there are path and strategies forward to mitigate them within a year's time, modulo focus and efforts from the collaboration.

Computational challenges will continue to remain an issue for the STAR experiment. We have already severely restricted the scope of data production campaign to <1.2 passes per year of data-taking, leaving no room for missteps or improvements. A further recent funding limitation (below flat budget for computing at our Tier-0 center) has made this even harder and driven STAR to seek to further expand its remote facility resources via the OSG software stack, essentially outsourcing temporary “boost” resources. STAR, pioneering Cloud technologies in the mid-2000s,⁵ is also seeking to evaluate the usability of a Docker-based approach on HPC clusters. A current pilot project at the NERSC/Cori ASCR facility shows promising results (95% WallTime/CPUtime efficiency), but much will be needed to make this (currently manually handled) work flow usable by the whole community. Within the context of the OSG, STAR would propose to make this happen in collaboration with NERSC.

2.2 Codes and Algorithms:

The STAR framework is a single-purpose plug-and-play framework allowing the insertion of processing modules (aka “Makers”) that are chained. It serves all the simulations and data-processing needs, while analysis may use the framework or plain ROOT-based analysis. The simulation is Geant3 or VMC/TGeo oriented. Most of our codes and algorithms are written in C++ with some inheritance from the FORTRAN era (Geant3). Our work flows also rely on access to calibration and condition databases. All codes are saved in a code repository system, and the end-user analysis is further accompanied by an “analysis note,” documentation and material describing in more details the analysis work flows. Our algorithms are based on standard Kalman Filter tracking, KNN, or Cellular Automata seed finding, Geant simulations and Markov Chain Monte Carlo techniques and data mining as well as MVA techniques at analysis levels.

3. Current and Future HPC Needs

3.1 Computational Hours:

The data processing of RAW data and reduction to DST are currently being done on approximately 15,000-core processing farms for production campaign cycles of one pass within a year. Additional resources at the 20% level are used for covering for the simulation and additional processing speedup demands (those resources come from a combination of a NERSC allocation and outsourced resources accessed via Grid). With farm's uptimes of ~ 95% and work flow efficiencies in excess of 95% (98% for raw data production), we estimate the demand to be ~118 Million CPU hours at the Tier 0 center and an additional 24 Million from remote facilities. User analysis is made over at least ~ 5,000 CPU cores and represents an additional ~ 25 Million CPU hours.

⁵ See for example J. Lauret, M. Walker, S. Goasguen et al. “When STAR meets the Clouds – Virtualization & 'Cloudified' grids experience”, [CHEP 2010 proceedings](#), *J. Phys. Conf. Ser.* **331** 062016 [doi:10.1088/1742-6596/331/6/062016](https://doi.org/10.1088/1742-6596/331/6/062016) (2011)

3.2 Parallelism:

Apart from the online HLT/L4 processing using Xeon/Phi, there is little room for the use of parallelism offline. While parallelization of the track reconstruction processing (60% of the work flow at best) could be enabled, it is unclear that the overall cost/performance would benefit considering the 40% serialized work flow. In addition to such fine-grained parallelism, our framework is ready for “Chain” and “Maker” coarse-grained parallelization.⁶ Processing has mostly been carried in a pleasantly parallel mode. Efforts will continue toward vectorization, but, unless cost/performance is at the sweet spot, we do not immediately see the further parallelization as part of our highest task priority.

3.3 Memory:

Most of our work flows require at least 2 GB of memory per core with some peak at 3 GB/core. However, with the introduction of denser detectors and potentially more complex events, we may expect and require 4 GB/core in outer years (beyond 2020). Shared memory as well as cache effects have not yet been studied in details in STAR, but the introduction of tools making this task much easier would be most welcome.

3.4 Scratch Data and I/O:

Offline reconstruction over 15,000 CPU cores could produce as much as 15 TB/day of derived data (+/- 1 TB/day), leading to a relatively modest offline requirement of ~1 GB/core per day on average. The input demand is about 2 times this amount. Each processing node, however, requires an average 15 GB of local store for moving an input file in, processing and generating the output locally, and moving it to Mass Storage. User analysis has been monitored to require an aggregate of about 12 GB/sec writes and 15 GB/sec reads over ~5,000 CPU core for output (hence 2.4 MB/sec writes and 3 MB/sec reads on average) but with a broad distribution. The most dramatic fast-detector analysis was showed to require up to 10 MB/sec read.

3.5 Long-term and Shared Online Data:

Our raw data storage needs are of the order of 10 PB/year but will decrease to a couple of PB for the two years of the BES II program and back to the higher standard need or less beyond 2020. Derived data sets are a factor of ~2 less, with 8 PB of MuDST kept “live” for user analysis in the Scalla/Xrootd disk space, insufficient to keep all years’ data. The introduction of the picoDST format would make it possible again to provide at least 5 years’ worth of data sets simultaneously. Currently, the data available “live” is accessible and shared among all users. Sample data sets or specific years are transferred over the WAN (using Grid tools) to Tier-1 or -2. Grid-based strategies to share data off Tier-0 are essential to our computing model.

3.6 Archival Data Storage:

The total space usage to date is 24 PB of RAW data (real and simulated) and ~9 PB of derived data. All of our RAW data are imagined to be kept forever, while only the vetted copy of a year’s derived data sets will be kept. At a conservative projected growth forward of 8 PB for 2 year for Run 17 and 18 and 2 PB for the BES II years, the RAW is seen as totaling $24+10+5+1+1 \sim 40$ PB and $9(+8)+6+3+1+1 \sim 30$ PB of RAW and DST storage, respectively, by 2020 (DST will see additional storage growth due to the ongoing past years’ data productions indicated in parenthesis). Beyond 2020, it is seen that only a fraction of the MuDST would be preserved, decreasing the DST storage demand. By 2025, a possible accumulation of

⁶ Fine, V., Lauret, J.; Perevozchikov, V. “Evolution of the STAR Framework OO model for Multi-Core era”, [ACAT08 proceedings](#), PoS ACAT08 040 (2008).

up to 70 PB and 35 PB for respectively RAW and DST are possible, but this projection is not clear as it will depend on the running scenario.

3.7 Workflows:

Our work flow was described in section 2.1. We intend to continue the same approach with a strong possibility to leverage resources such as the one provided by NERSC/Cori, which we expect to be suitable for our needs. The use of Grid or Cloud resources will continue to be part of STAR's standard way of processing. A lower increase rate of resources would have major impact on our ability to deliver science in a timely manner and any shortage at the Tier-0 facility would need to be found as borrowed resources elsewhere.

3.8 Many-Core and/or GPU Readiness:

As previously noted, it is unclear if the use of parallelization in our framework would provide an overall benefit. Similarly, the use of thread in a many-core context is doubtful but will be reassessed as novel hardware becomes available. The use of hardware accelerator would be studied for simulations, random number generators and the like, and the available GPU-aware ROOT functionalities. Several parts of our framework are thread, multi-core, and/or GPU ready.

3.9 Software Applications, Libraries, and Tools:

Typical third-party software for STAR includes the CERN/ROOT package, Geant (3 or 4), and the necessary MySQL and XML2 library. We also require FORTRAN, C, and C++ compilers, although by 2020, FORTRAN's needs *may* have phased out entirely. In the past, our tested work flows on the Grid and Cloud using virtualization has carried along the necessary packaged software components including 3rd parties, but the compilers and MySQL/XML2 libraries (installed in VMs, tailored to the embedded and not installed by STAR). Our test work flow on Cori follows the same approach.

3.10 HPC Services:

The Grid-based jobs submission approach requires the maintenance of at least a Grid gateway for submission and at least one separate and a data transfer node connected to a central storage. Grid security is also most needed and provided by the OSG software stack. At NERSC, an ASCR facility, STAR highly benefits from having redundant MySQL servers. An automated scalable architecture relying on virtualization and load detection would be a nice add-on, the next natural frontier being "service on-demand." STAR is supportive of a single sign-on federated authentication service that is not orthogonal to the Grid paradigm.

3.11 Additional Needs:

STAR computing research for the past many years included the study of advanced techniques for efficient multi-site data transfer.⁷ From AI to Constraint Programming to Network Flow Maximization, all techniques show a large benefit in coordinated job submission steered by an intelligent planner. Up to a 30% speedup on global work flows has been demonstrated, a significant gain over the harvesting of

⁷ See, for example, [Michal Zerola, Roman Barták, Jérôme Lauret, and Michal Šumbera. "Using Constraint Programming to Plan Efficient Data Movement on the Grid". In *Proceedings of the 21st IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 729–733, [IEEE Computer Society, 2009](#)] [Michal Zerola, Jérôme Lauret, Roman Barták and Michal Šumbera – "One click dataset transfer: toward efficient coupling of distributed storage resources and CPUs", ACAT 2011 proceedings, *J. Phys.: Conf. Ser.* **368** 012022 [doi:10.1088/1742-6596/368/1/012022](#) (2012)] and [Dzmitry Makatun, Jérôme Lauret, Hana Rudová and Michal Šumbera, "Planning for distributed workflows: constraint-based coscheduling of computational jobs and data placement in distributed environments", ACAT2014 proceedings, *J. Phys.: Conf. Ser.* **608** 012028 [doi:10.1088/1742-6596/608/1/012028A](#) (2015)]

global resources and one that cannot be left under-considered. Lack of support for integration and maintenance activities has hindered the propagation of knowledge across the community. More important than training, programs supporting moving research endeavors of this nature to practical application for the benefit of the entire community should be created for well-justified activities. The push by agencies to “preserve” data would imply, in our views, the need to transfer a copy of our critical data sets “elsewhere” for maximum safety. However, to date, we have not been able to settle an agreement that would achieve just that, the long-term preservation of not being in scope with the ASCR facilities’ model.

Requirements Summary Worksheet

Code: STAR ROOT	Column 1: CurrentUsage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	170 M	x2	x2
Computational node hours (homogeneous many-core)	All of the above	-	-
Computational node hours (with GPU or accelerator)	N/A	N/A	N/A
Memory per node	2 GB/core	3 GB/core	4 GB/core
Aggregate memory	? TB	? TB	? TB
Data read and written per run	RAW 10 PB DST 6 PB	x0.4 x0.5	x1 x0.2
Maximum I/O bandwidth needed	RAW 1.6 GB/sec DST 15 TB/day User 15 GB/sec	x1 x2 x2	x1 x2 x1?
Percentage of runtime for I/O	RAW N/A DST 5% User 10%	-	-
Scratch file system space needed	RAW 100 TB DST 15 GB / “job”	x1	x1
Permanent online data storage	RAW N/A DST 8 PB	- 20 PB	- 20 PB
Archival data storage needed	RAW 24 PB DST 9 PB	RAW 40 PB DST 30 PB	RAW 70 PB DST 35 PB

Precision Extraction of Quark-Gluon-Plasma Properties

Steffen A. Bass (Duke University)

1. Description of Research

1.1 Overview and Context:

For about 10 microseconds after its creation in the Big Bang, the universe was in a state called the Quark Gluon Plasma (QGP) in which quarks and gluons, the basic constituents of the strong nuclear force, Quantum Chromo Dynamics (QCD), roamed freely. Due to the universe's rapid expansion, this plasma cooled and went through a phase transition to form hadrons—most importantly, nucleons—that constitute the building blocks of matter as we know it today. The Relativistic Heavy-Ion Collider (RHIC) at Brookhaven National Laboratory was specifically built to observe and study the QGP phase of matter, and similar studies at the CERN Large Hadron Collider (LHC) have recently confirmed the QGP discovery.

The central problem in studying the QGP is that its lifetime is so short that only the ashes of its decay (in the form of hadrons) can be detected. By comparing measurements taken at the Collider facilities to computational models that simulate the nuclear collisions in which the QGP is created, we seek to establish unambiguous connections between the transient (quark–gluon) plasma state and the experimentally observable signals. A single model-to-data analysis can require in excess of 10^6 CPU hours and is commonly run on grid computing facilities, requiring sufficient network throughput to collect and analyze the output generated on individual nodes.

1.2 Research Objectives for the Next Decade:

The objective for the next decade is a comprehensive quantitative determination of the properties of the QGP, in particular, its initial state and the temperature dependence of its transport coefficients, with quantified uncertainties. The computational modeling environment will rely on a suite of codes currently being developed that are capable of modeling the entire history of a relativistic heavy-ion collision from the onset of the collision through QGP formation, the confinement phase transition to hadronic final state interactions and freeze-out.

Ideally these research objectives would be met as a community effort in the context of a collaboration of all US nuclear theory groups involved in the modeling of and large-scale knowledge extraction from relativistic heavy-ion collisions. Currently there are 4-6 institutions in the US that would be interested in pursuing these objectives, possibly leveraging the ASCR computing facilities and resources through a seamless interface for job submission (The Open Science Grid, OSG, software stack has successfully been used⁸).

2. Computational and Data Strategies

2.1 Approach:

State-of-the-art computational models are fairly complex and rely on more than a dozen parameters that encode the physical properties of the system we wish to extract. Scanning the full parameter space at a fine resolution is computationally prohibitively expensive (10^8 CPU hours for a single analysis). We

⁸ Nuclear physics and computer science meet on the OSG, [Open Science Grid news report, Feb. 24 2014](#).

therefore integrate Bayesian inference and model surrogate algorithms⁹ into our simulation environment that hugely improve the efficiency of exploring the multi-parameter model space and allow for the determination of high-likelihood parameter ranges by comparison with the experimental Petascale data sets. Currently we utilize the OSG to provide us with the necessary compute cycles. The OSG offers the flexibility of providing these cycles “on demand” when we are ready to run an analysis with minimal organizational overhead (important in a University environment with changing personnel at different levels of training as well as a very dynamic research program). The main challenge we face with the OSG is the saturation of available resources. Anticipating future requirements at 10 to 50 times the current resources will be very challenging to obtain from the OSG in its current state.

2.2 Codes and Algorithms:

The physics model is based on relativistic viscous fluid dynamics coupled to Boltzmann transport theory in multi-component systems, radiation transport theory, and source imaging techniques. Model surrogate algorithms are based on Gaussian process emulators, and the determination of high-likelihood parameter areas utilizes Markov Chain Monte Carlo techniques. To generate the surrogate model training data, the physics model is run for an ensemble of parameter values determined by a Latin hypercube, and the individual calculations are executed on the OSG.

3. Current and Future HPC Needs

3.1 Computational Hours:

One model-to-data study for the extraction of temperature-dependent transport coefficients using our current 2+1D vRFD code requires about 21 million CPU hours (100K events per parameter point, 500 design points for 12 parameters, three beam energies, 4-5 centrality bins and 0.15 hours per event). Subject to availability of resources, our group could run up to four such studies per year, so current annual usage is on the order of 100 million CPU hours. Within the next five years, we hope to develop a new 3+1D viscous RFD code that runs on GPUs and anticipate a speedup by two orders of magnitude (compared to a 3+1D vRFD single CPU code) if GPUs become available on a large scale. The larger the computational throughput and the quicker the individual studies can be completed, the larger the number of such studies that can be accomplished in a given year, since the outcome of one study/analysis informs the goals for the next one. With increasing computational capacity, we anticipate a doubling of our needs by 2020 and another doubling by 2025.

3.2 Parallelism:

Our current work flow and codes are not parallelized and are not well suited for standard parallel computing techniques. They run exceedingly well on individual CPU’s and are very well suited for standard Grid computing applications (of course they could make use of multiple cores independently as our calculations are typical examples of embarrassingly parallel workflows). We are in the process of developing a new 3+1D vRFD code with GPU acceleration that would become the standard within the next five years. We anticipate a speedup of at least a factor of 100 with the new code (compared to single CPU execution). Similar, albeit less dramatic speed-ups, could also be obtained using standard parallelization techniques.

⁹ Applying Bayesian parameter estimation to relativistic heavy-ion collisions: simultaneous characterization of the initial state and quark-gluon plasma medium - Jonah E. Bernhard, J. Scott Moreland, Steffen A. Bass, Jia Liu, Ulrich Heinz, [nuch-th arXiv:1605.03954](https://arxiv.org/abs/1605.03954) preprint.

Training of scientific personnel in modern computational techniques is presently a bottleneck as there is little funding available for the computational development of physics models and the training of physics students in modern computational techniques. We are currently leveraging NSF funding through the Software Infrastructure for Sustained Innovation (SSI) program for the development of the GPU based vRFD code. Focusing on GPU acceleration vs. standard parallelization is mostly due to availability of such targeted funding opportunities.

3.3 Memory:

Current 3+1D vRFD codes require approximately 2–4 GB of memory on a single CPU.

3.4 Scratch Data and I/O:

We require about 200–400 GB per CPU scratch storage for running the computational model and about 200 TB scratch storage for running the data analysis portion of the computation. The data analysis portion only takes place $O(500)$ times over the course of the full computational campaign. I/O during the running of the computational model requires very little time, less than 5% of running time is devoted to I/O; therefore bandwidth of I/O is noncritical.

3.5 Long-term and Shared Online Data:

Each computational campaign (of which we anticipate currently 4 per year with growth to 8 in 2020 and 16 in 2025) requires about 200 TB of space. We only need to keep data from a few campaigns active at any time, therefore approximately 800 TB (current) with modest growth to 1,000 TB (2020) and 2,000 TB (2025) should be sufficient. Data access occurs by all members of the research collaboration, i.e. from 3–5 institutions in the US. While we would like to make all data publicly available, we do not anticipate significant demand for access.

3.6 Archival Data Storage:

Data archival needs will be mostly in response to funding agency requirements. We anticipate archival needs of about 800 TB per year if we choose to archive all data generated during our analysis. The amount of long-term and archival data can be dramatically reduced at the expense of re-running the computational model and thus reproducing previously generated data. Access delay for archived data is not an issue.

3.7 Workflows:

Our current work flow is centered on the OSG as main source of our computing cycles. Our computational model consists of a sequence of several codes (initial condition code, vRFD, “particleization” code, Boltzmann transport, and event analysis) that are being handled by a steering Python script embedded in a Condor job description for OSG scheduling and operations. Each OSG job is tuned to run approximately 10 hours on a single OSG compute node, and the model output of that job is processed on the OSG node before the processed data is brought back via grid-ftp and written to our local analysis cluster. The details of the work flow depend on the per-event execution time of the computational model and the amount of data storage space available on the compute node as well as on the analysis cluster. In the future we anticipate adapting the work flow to the computational and storage resources at hand.

Ideally we would like to significantly reduce the total time for a given analysis from several months to a few days, since such a short time scale would allow for an informed and near interactive exploration of model features and data sensitivities that is not feasible under the current time scales.

3.8 Many-Core and/or GPU Readiness:

We are in the process of developing a new 3+1D vRFD code with GPU acceleration that would become the standard within the next five years. We anticipate a speedup of at least a factor of 100 with the new code (compared to single CPU execution). That speedup would translate to a significant reduction of required node hours for the same level data and scientific output or, at the same expenditure of computational resources, would allow for a vast increase in the scientific capabilities of our analysis (albeit with a factor 100 more data to be analyzed, stored, and archived). We would be able to significantly increase the number of experimental measurements that would be included in our analysis, and this would allow for much more precise constraints on a larger number of QGP transport coefficients.

3.9 Software Applications, Libraries, and Tools:

Our current computational setup requires FORTRAN, C, and C++ compilers as well as the SciPy ecosystem. Additional libraries include HDF5, H5py, GSL, and boost. For our GPU-based vRFD development, we will require OpenCL and CUDA.

3.10 HPC Services:

We would be very interested in exploring data analytics as well as advanced visualization techniques and collaborative tools for improving our work flow and data representation.

3.11 Additional Needs:

N/A

Requirements Summary Worksheet

Code:	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	100M	2x	4x
Computational node hours (homogeneous many-core)			
Computational node hours (with GPU or accelerator)**	20,000,000 (currently in development)	4x	8x
Memory per node	4 GB	4 GB	4 GB
Aggregate memory	TB	TB	TB
Data read and written per run	0.2 TB	0.4 TB	1 TB
Maximum I/O bandwidth needed	10 GB/sec	20 GB/sec	40 GB/sec
Percentage of runtime for I/O	5%	5%	5%
Scratch file system space needed	100 TB	200 TB	400 TB
Permanent online data storage	800 TB	1,000 TB	2,000 TB
Archival data storage needed	800 TB	3,200 TB	8,000 TB

*Core hours for conventional processors.

**Node hours for GPU or accelerator usage.

FRIB/NSCL Active Target Time Projection Chamber (AT-TPC)

Lead Authors: Wolfgang Mittig and Yassid Ayyad (ayyadlim@nscl.msu.edu)

(National Superconducting Cyclotron Laboratory, Michigan State University)

1. Description of Research

1.1 Overview and Context:

The Active Target Time Projection Chamber (AT-TPC) is a novel active target well adapted for nuclear reactions with exotic beams at low bombarding energies. The AT-TPC uses a readout technology based on a micromegas gas amplifier that consists of a micro mesh coupled to a pad plane segmented in 10,240 pads. The GET electronics (General Electronics for TPCs) allows the sampling of the pulse induced in each pad with a variable sample rate with a maximum of 100 MHz. These conditions lead to a large amount of data (of the order of TB per experiment) that will depend on the trigger rate at which the experiment runs. From the pulse shape analysis of each one of the registered pulses, one can extract a three dimensional hit pattern depicting the reaction inside the detector. Contrary to other typical nuclear physics experiments, inferring the physical observables from the reaction requires the use of sophisticated techniques and algorithms with an intensive use of computing resources.

1.2 Research Objectives for the Next Decade:

The AT-TPC is a versatile detector that can be used either with high- or low-energy beams. The project has a wide scientific program covering reactions of astrophysical interest, cluster structures in exotic nuclei or nuclear spectroscopy through direct reactions, among many others. With the advent of the Facility for Rare Isotope Beams (FRIB), with unprecedented intensities, it is capital to develop reliable hardware/software tools to efficiently distribute, store, and analyze the data acquired with the TPC in order to validate the setup condition (maximize beam time) and discriminate rare events which will be one of the main goals when using very exotic radioactive beams.

2. Computational and Data Strategies

2.1 Approach:

One of the main problems is how to efficiently validate the data in the fastest possible way. The AT-TPC has a distributed data acquisition system that generates a certain number of files that must be merged to access into an event-by-event reconstruction mode. An efficient unpacking routine would include a parallelized code that deals with these multiple files at the same time. The same paradigm can be applied to the next steps of the data analysis. A great deal of effort is currently being carried out within the world of parallel computing at different levels. The AT-TPC software project envisions how to benefit from such future developments by incorporating the most widely used libraries/frameworks for parallel computing.

2.2 Codes and Algorithms:

Our algorithms are designed in a highly modular way. Our code is written mostly using C++, C, Python, and FORTRAN. Different tasks, such as unpacking, sorting, and analysis, can be independently modified and acceded in a hierarchical way. We developed several pulse-shape analysis routines as well as several pattern recognition algorithms. For the analysis we use typical Monte Carlo routines and also Kalman filters in the future. We aim to fully parallelize our code using OpenMP, OpenMPI, and CUDA.

3. Current and Future HPC Needs

3.1 Computational Hours:

In 2016 we have been testing our reconstruction routines parallelized with the CUDA framework on a Jetson TX1 embedded system with 256 NVIDIA CUDA cores and also on the HPC Cluster at MSU. In particular, we tested the speed performance for heavy floating point calculations, where the GPU is expected to provide a better performance compared to multithread CPU calculations. Our preliminary tests show an increase in speed by a factor of the order of 5–10 in such calculations. This will significantly reduce the critical number of computational hours needed to analyze data taken with the AT-TPC. This improvement will be needed in order to counterbalance the massive amount of data from future FRIB experiments, with beam intensities 10^3 – 10^6 times larger. In total, we currently use 10 threads (cores) to unpack the data in parallel. Then, our aim is to use an event-by-event analysis structure, where each core takes the analysis of a single event, and this can be scaled to the number of available cores (for example around 28 in the current HPCC), making the analysis acceptably fast. The last step would be to speed up all the calculations using CUDA cores.

3.2 Parallelism:

The unpacking process uses `std::thread` to create single threads for each one of the raw data files. For the analysis, we currently use OpenMP library, provided with the most modern GCC versions. For the most demanding floating-point calculations we are implementing CUDA parallelization.

3.3 Memory:

Note that future systems may have much less memory as a function of peak performance than systems have today. Note also that the memory hierarchy will be potentially complex (on-chip fast memory and significantly slower to off-chip memory).

The AT-TPC DAQ is a stand-alone system so we will not comment on its memory requirements. On the other hand, for the data analysis we do not have a current online software that could greatly benefit from a shared memory system. We plan to develop an online framework in the future, but the requirements are not clear yet since the amount of streamed data sent from our DAQ may vary. Our full unpacking-analysis task requires an amount of ram memory of the order of 2 GB for 1 run (around 1 hour).

3.4 Scratch Data and I/O:

The size needed for running an experiment would be of the order of 1 TB/day, having a bandwidth of the order of 10–100 MB/s depending on the trigger and electronics configuration. Nevertheless, we perform a data reduction by only recording the hit pattern for each event and several check parameters.

3.5 Long-term and Shared Online Data:

The data streaming during an experiment is of the order of hundreds of MB/sec (300–500 MB/sec) at a beam rate of 100 Hz with the average settings for the electronics (which may vary substantially). The totality of the raw data is stored in different drives, and we perform a substantial reduction of the data to be used as online data. This data reduction is of the order of a factor of 5, with the reduced data of the order of 1 TB. Nonetheless, it is also desirable to keep the raw data when testing and debugging the detector. In the future, this total amount of online data would range from 10 to 100 TB. A comfortable

and optimum solution would be to send the online data to a HPC cluster in order to provide a shared access to the users in order to analyze the data.

3.6 Archival Data Storage:

Since we perform a data reduction, the amount of data is currently reduced to few hundreds of GB per experiment, having in total three experiments performed with the AT-TPC.

3.7 Workflows:

Each user/developer can work in any of the different tasks in which the framework is divided without any interference. Versions are managed through git and github.

3.8 Many-Core and/or GPU Readiness:

Our framework is ready to use the most novel parallel computing technologies. We plan to use HPCC resources as well as portable stand-alone systems for parallel computing such as nVidia Jetson TX1 (we are testing the system).

3.9 Software Applications, Libraries, and Tools:

We use the most common scientific libraries for C++ and python, such as ROOT, Geant4, and SciPy.

3.10 HPC Services:

Unknown at this time.

3.11 Additional Needs: Not addressed here.

Requirements Summary Worksheet

Code: FRIB/NSCL at-TPC	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	NA	NA	NA
Computational node hours (Homogeneous many-core)*	180 per expt.		
Computational node hours (with GPU or accelerator)	NA	NA	NA
Memory per node	2 GB	2+ GB	2+ GB
Aggregate memory	NA	TB	TB
Data read and written per run (experiment)	0.02 TB per run	TB	TB
Maximum I/O bandwidth needed	0.01–0.1 GB/sec	GB/sec	GB/sec
Percentge of runtime for I/O			
Scratch file system space needed	10 TB per expt.	TB	TB
Permanent online data storage	NA	TB	TB
Archival data storage needed (cumulative for all experiments)	1 TB per expt.	TB	TB

* Node hours for homogenous many-core architectures.

Data Acquisition and Online Analysis at FRIB

**Lead Author(s): FRIB Data Acquisition Working Group; Robert Varner
(varnerri@ornl.gov)**

1. Description of Research

1.1 Overview and Context

The FRIB Data Acquisition Working Group does not have a direct research agenda, but instead works to understand and implement effective data acquisition tools for FRIB users. FRIB data acquisition differs from that at the Relativistic Heavy-Ion Collider (RHIC) or the Continuous Electron Beam Accelerator Facility (CEBAF). Experiments will often be set up and run to completion in one to three weeks, without a long time to debug and optimize acquisition, control, or analysis software. FRIB experiments will need to use larger detector systems with many more channels to effectively utilize the beams of rare isotopes for the wide variety of nuclear physics questions that can be posed. The architecture of the data acquisition system must enable users, often not highly trained programmers, to collect data efficiently in the short time of the setup and running. In addition, we are in the midst of a paradigm shift in nuclear physics data acquisition, from the use of integrating digitizers with hardware waveform shaping and delay lines, to signal digitization, and eventually the storage of the raw detector waveforms for online and offline analysis. The dramatic increase in data volume will pose a significant challenge to all aspects of our computing environment.

1.2 Research Objectives for the Next Decade:

The scientific goals of FRIB are well known: enabling studies of nuclear structure and reactions closer to the limits of nuclear stability to inform models of nuclei and the behavior of nuclei in extreme astrophysical systems. The FRIB data acquisition working group will enable experimenters to acquire data efficiently in the short time available for experiment setup, to effectively use larger ensembles of detectors, to prepare for fully digital data acquisition and finally, to be able to analyze all the data in real time, and to make well-informed decisions about the experimental design and operation. The change to fully digital data acquisition will lead to a data volume increase of two to three orders of magnitude at the front end, demanding a comparable increase in processing power. The need for faster online data processing will require an additional two orders of magnitude or more increase in processing power. All of these need a networking capacity sufficient to handle the data. By 2025, we expect a large fraction of data acquisition to use “digital data acquisition” techniques.

2. Computational and Data Strategies

2.1 Approach:

Low-energy nuclear physics data acquisition collects data from digitizing modules in instrumentation backplanes that support readout speeds from tens to hundreds of MB/s that respond to randomly occurring events, move the data with a minimum of dead time to storage, and analyze the data in real time to diagnose the operation of an experiment. Current data acquisition uses a few general-purpose compute cores for readout, data sorting, and storage for data rates of 10 MB/s or so, and a network of 1Gb/s for a net data rate of 32 TB per year at a facility. The data acquired are eventually time-ordered, and to the extent that blocks of data in events are independent, the data are amenable to trivially parallel processing. Although existing systems often utilize multiple threads or processes to perform

these tasks, the algorithms are not explicitly parallel and are not designed to take advantage of large numbers of cores. Data analysis can often significantly lag behind the experimental data acquisition and data collection, and movement must minimize the dead time imposed on the readout.

There are three changes coming in the next decade. Experimenters are interested in faster event processing, to allow them to examine all the data faster than the acquisition rate; detector arrays will grow significantly to more effectively use the rare beam; and experimenters are changing to waveform recording rather than hardware integrating digitizers. Addressing these issues will require more parallel data paths to handle increased data movement, taking advantage of the trivial parallelism of the event data. More processor cores and automation in the framework will distribute the data to the cores. We will need to provide significantly more storage for event data for post-experiment analysis of the data.

2.2 Codes and Algorithms:

Data readout, time ordering, transport, and storage are standardized in professionally designed frameworks that are customized by experimenters to include the experiment-specific readout and analysis. The algorithms are simple serial programs in languages ranging from FORTRAN to C++ or JAVA, usually, created by researchers. Currently, most hardware digitizes analog signals that are read out in blocks of data ordered in buffers for each module. Experiments produce data at a rate of a few bytes per channel per signal. The data read from several instances of front-end hardware are transported to an event builder, which sorts the buffers together into time correlated “events.” Each event captures single-beam particle interactions in the target. The correlated event stream is then transported to file storage, written to file servers, and presented to online analysis. The online analysis transforms the event signals into physics related parameters and generates relevant histograms.

If the readout is digital, the waveforms of the signals are digitized by flash analog-to-digital converters. In that case, the waveforms are transported. The data volume is 10 to 100 times greater than for analog readout, and each waveform is tagged with a clock to help in later event ordering. The sorting and storage processes are similar, but deal with the substantially larger volume of data. The final steps of analysis are similar, but must now be analyze the features of each waveform to extract physically meaningful parameters.

3. Current and Future HPC Needs

Most of our estimates below are based on an optimistic assumption that FRIB and its predecessor will run for 5,000 hours per year. Analysis will, most likely, run full time for 8,000 hours per year.

3.1 Computational Hours:

Current requirements vary significantly. On average, readout and data storage take between two and eight cores while experiments run, for a total of 40k core hours/year. The online analysis, usually single-threaded, takes place on one or two CPU cores, one devoted to data transport and the other to online analysis, 200 to 400 core hours per week. If waveform digitization dominates, this could require 10 to 1,000 times the processing power of existing analog systems. Additionally, the growing complexity of experimental setups requires additional increases in online processing power, to allow examining the data in faster than real time, perhaps a 100 to 10,000 times total increase based on discussions with groups. Waveform analysis requires more processing, at least comparable to the scale of the data increase, if not more.

3.2 Parallelism:

Current data acquisition systems exploit multiple lightly synchronized processes and multiple nodes where the equipment is easily isolated. The desire for faster-than-real-time replay of the data will require analysis code that can explicitly utilize parallelization on a much larger scale. When waveform readout begins to dominate, the data acquisition systems will require multiple nodes to serve sufficient bandwidth to the hardware and network to move the data. Highly parallel systems including GPUs will perform the required waveform analysis.

3.3 Memory:

Current memory usage is small, 32 MB per node, with analysis requiring 256 MB just for data buffer management. Increasing the parallelism of the analysis will require much more sharing of the data stream and synchronization mechanisms to preserve the data stream. We expect the memory per node to increase and the aggregate to rise, with the larger number of processors.

3.4 Scratch Data and I/O:

Scratch data are rarely used now, because of the large bandwidth penalty of going to disk. The very largest analyses may require scratch I/O in the future. High-speed local scratch storage will be important for large cluster access to the most recent runs, to reduce I/O bottlenecks in the analysis. The computational load of data acquisition is usually I/O bound. Currently it is limited by the 10 MB/s of standard experiments and by the need to create large numbers of large histograms and other graphical representations of the data. The analysis may need to read back the event stream from mass storage, more than doubling the required bandwidth. With larger detectors and waveform analysis, the I/O bandwidth requirements for analysis will grow without bound.

We intend that our analysis system will be I/O bound, at which point we need to be able to read the data faster than it was acquired, which might be at GB/s.

3.5 Long-term and Shared Online Data:

The National Superconducting Cyclotron Laboratory (NSCL) currently stores data from an experiment for a relatively short time, transferring it to the researcher on tape. Thus, they have 20 TB available for a yearly load that must approach 50 TB per year. The increasing size of the data sets, the federal requirements for data preservation, and the increasing size of collaborations argue for a substantially larger file store at FRIB, along with computational resources to reduce waveform data to an easily transferable physics data set. A file server of 150 TB should be a minimum today, hosting all experiments for three years. This will grow with the size of the raw data stream in data density and number of channels. By 2025, this will be many tens of petabytes, particularly if device readout should evolve to much higher data sample rates, such as 1 G sample/s.

3.6 Archival Data Storage:

The requirements for archival storage will be driven by the large amount of data from experiments with many channels, high rate sampling of waveforms, and high count rates. A maximum case estimate for one experimental device would produce 1.3 PB/week, probably running for two weeks per year at FRIB. We estimate, based on 3.5, that a 100-PB scale archival storage is required.

3.7 Workflows:

The work flow of data acquisition was described in section 2.2. This work flow will change somewhat over time. Over the next decade we will move data acquisition to depend less on custom electronics solutions for each detector system to utilizing more the networking, parallel computing, and high-

density storage solutions driven by the larger technical marketplace. By 2025, most experiments will digitize waveforms and process them in the analysis. The work flow will still be transferring waveforms with time stamps from hardware to intermediate processors that time-order the blocks of data coming from many cores performing the readout. This time-ordered stream will move to intermediate data handling that distributes the data to a long-term storage device and a highly parallel network of computers. The parallel computers will analyze the waveforms to create a new stream of time-ordered physics parameters that will be saved to storage and distributed to event analysis processors that will create quantitative graphical presentations of the event data. All this must be accomplished in less time than it takes to run the experiment by 2025.

3.8 Many-Core and/or GPU Readiness:

Our codes are not ready for many-core and GPU systems, but our data has a trivial parallel structure (discrete events) that can be easily exploited by such systems. GPUs are useful for many of the problems in waveform analysis.

3.9 Software Applications, Libraries, and Tools:

All our software is developed by the FRIB community and generally requires little external code except for compatibility with CERN ROOT and related codes, such as CLHEP.

3.10 HPC Services:

3.11 Additional Needs:

High-performance and reliable interconnects between the FRIB front end and storage and the computing resources. The computing resources need to be reserved for real-time analysis during experiments.

Requirements Summary Worksheet

Code: FRIB Data Acquisition	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	30,000	100X	1000X
Computational node hours (homogeneous many-core)	0	**	++++
Computational node hours (with GPU or accelerator)	0	**	++++
Memory per node	16 GB	64 GB	256 GB
Aggregate memory	0.05 TB	5 TB	50 TB
Data read and written per run	2.0 TB	100 TB	1000 TB
Maximum I/O bandwidth needed	0.050 GB/sec	1.0 GB/sec	10.0 GB/sec
Percentage of runtime for I/O	50%	50%	50%
Scratch file system space needed	0.0 TB	200.0 TB	2000.0 TB
Permanent online data storage	20.0 TB	10,000 TB	100,000 TB
Archival data storage needed	50.0 TB	10,000 TB	100,000 TB

* Core hours for conventional processors.

** We intend to explore these technologies for enhancing the throughput of our analysis. We cannot estimate yet our effectiveness at this.

High-Performance Computing for FRIB Design, Operation and Future Upgrades

Lead Author(s): Dali Georgobiani, Mauricio Portillo, Marc Hausmann, Mikhail Kostin (kostin@frib.msu.edu), Georg Bollen, and Frederique Pellemoine (Facility for Rare Isotope Beams, Michigan State University)

1. Description of Research

1.1 Overview and Context:

The Facility for Rare Isotope Beams (FRIB) is an accelerator facility under construction at Michigan State University. The facility will utilize a broad range of primary ion beams from 16-O to 238-U with a beam power of up to 400 kW and energy of 200 MeV/nucleon for 238-U in its baseline configuration to produce rare isotopes. A possible facility upgrade will include an increase of the beam energy up to 400 MeV/nucleon for 238-U and addition of new light ion beams down to 3-He and protons for ISOL operations.

High-performance computing (HPC) techniques have been in use to address various aspects of the FRIB facility design, such as the bulk radiation shielding, beam optics, and individual beam line elements. We also expect to use HPC to support facility commissioning, operation, and possible upgrades. HPC is mainly required to obtain statistically significant calculation results due to the size, complexity, and multitude of the models.

1.2 Research Objectives for the Next Decade:

Most of our calculations related to the design of the baseline configuration of FRIB have been completed. The facility is scheduled to have its first beam in 2020-2022. We expect to reach the maximum beam intensity within the first five years of operations. Most of our near-term effort is expected to be directed toward the support of the facility commissioning and, later, toward the operation into the next decades. It is also possible that we will have to address potential upgrades to the facility around 2025, after the commissioning will have started. The upgrades may include the increase of the primary beam energy and the design of an Isotope Separation Online (ISOL) target station. The experience of the FRIB radiation transport group can be leveraged to provide support for other national or international projects involving high-power accelerators.

2. Computational and Data Strategies

2.1 Approach:

One can separate the software being used into two categories with fundamentally different structure. One set of software and tools utilizes calculations based on a number of individually independent events. All our radiation transport and beam codes (MARS15, PHITS, MCNPX/6, COSY, LISE++), in principle, employ this approach in one way or another. This inherent parallelism allows running our calculations using either parallel computing (with Message Passing Interface (MPI)) or, alternatively, as a set of single-CPU jobs with consequent data analysis. All our calculations are "loosely coupled," which makes it possible to use computer farms with relatively low bandwidth. The criterion for successful calculations is acceptable statistical uncertainties. Given the sheer size and complexity of our model

(practically the entire FRIB facility), obtaining the required precision is often a challenging task, even though we use various variance reduction and biasing techniques. These calculations are also performed for many primary beams and many (in thousands) secondary beams.

The second set is the ANSYS code and its various modules. It is used for thermal and stress calculations for various components in the FRIB accelerator and target complex. A typical single-CPU job can normally handle calculation for only one beam impact, so we have to use approximations. Multiple beam pulses and more realistic cooling (like dual-phase liquids) need more computer power. We are currently experimenting with a parallel version of ANSYS and plan to rely on it heavily in the future. These calculations will also require considerable memory size and disk space.

In principle, both sets of tools are well established, and ANSYS is a commercial code. We do not expect to further develop them, and we do not expect any changes in approach over the next decade.

2.2 Codes and Algorithms:

The codes MARS15, PHITS, and MCNPX/6 are general-purpose radiation transport codes. They calculate interactions of various particles and their transport through 3D models defined with materials and geometry. The codes are capable of calculating several radiation field characteristics, such as particle fluxes, radiation doses, energy deposition, residual activation, radionuclide production, and the radiation damage to materials, to name some. All these codes are inherently parallel; a typical calculation consists of many individually independent events. A large number of such events are typically needed to obtain a sufficient level of accuracy. Very often we use a multi-CPU version of the codes, but in some cases it becomes an inefficient approach. If a large amount of material is present (like many meters of the radiation shielding), it is more practical to use variance reduction techniques. In one of those techniques, the geometry is split in parts, and we save particles that cross the boundaries of these parts, so that we can re-use the saved particles later, cycling them many times (source split technique). In this case we run many single-CPU jobs and save the data on disk.

The beam transport and optics codes LISE++ and COSY are used to find optimal beam optics settings for specific rare isotope beams. Depending on the problem solved, there is also parallelism in the codes—each beam particle is tracked individually through the optimized beam line. The problem here is the number of possible primary and secondary beam types. For example, we expect several thousand possible secondary beams with various charge states to be transported through the FRIB pre-separators.

The ANSYS family of codes (ANSYS Workbench, ANSYS CFX, ANSYS LS DYNA, and others) is used for stress and thermal calculations. It has been heavily used for beam element design, especially for such critical ones as the production target and the beam dumps. The future analysis of the target and beam dumps foresees the multi-cycle time-dependent loads and two-phase fluid flow analysis. This will require essential computational resources.

3. Current and Future HPC Needs

3.1 Computational Hours:

Our estimate for the computational hours is based on the experience of the previous and current calculations conducted as part of the FRIB design effort, and also on expectations of the effort for the facility commissioning, operation, and possible upgrade.

The average current use of the radiation transport codes MARS15, PHITS, and MCNPX/6 is approximately 10,000 CPU hours per week (conventional), resulting in approximately 500,000 CPU hours/year. During the operational phase of the facility, we expect to use the codes to calculate the radiation fields and their impact on the beam elements and radiation exposure to personnel and public for every secondary beam, which can be changed every week or two (depending on the requested beam type by the physics program). In addition to that, we will be involved in other supporting calculations, and possible calculations needed to upgrade the facility to higher primary beam energies and the new ISOL target station. Since the turnover of the results is expected to be faster in the future than currently, the required estimates are 3,000,000 CPU hours/y by 2020 (possibly soon after the start of the facility operation), and 6,000,000 CPU-hours/y by 2025.

The radiation transport is often coupled with the use of beam optimization codes COSY and LISE++. The current use of COSY and LISE++ is 20,000 CPU hours/y, and the projected need is 50,000 CPU hours/y in 2020, and 220,000 CPU hours/y by 2025. Note that the scaling for COSY and LISE++ is different than for MARS15, PHITS and MCNPX/6 because the latter is used for more calculations than only beam optimization.

The current use of ANSYS is 20,000 CPU-hours/y. The expected need is 160,000 CPU-hours/y in 2020, and 320,000 CPU-hours/y in 2025. Multi-core CPUs are preferred.

3.2 Parallelism:

The parallel versions of the radiation transport codes MARS15, PHITS, and MCNPX/6 were tested on a variety of clusters and supercomputers. A significant slowdown in speedup is seen at a few thousand CPUs. The parallel versions typically use Message Passing Interface (MPI). The parallelism is coarse-grained (although MCNPX/6 allows the use of threads for more fine parallelism). The calculations are conducted as either multi-CPU jobs or a batch of many single-CPU jobs. The existing level of parallelism is sufficient for our applications, and we do not plan to increase it. The situation with COSY and LISE++ is similar. No GPU use is currently possible. LISE++ does not currently allow parallel computing, but an effort is under way to develop a parallel version of the code.

ANSYS is a commercial code, and we will not be able to modify it. ANSYS is available in multi-CPU format. We were able to use up to 16 cores. We would like to extend it further to at least 32 cores or more.

3.3 Memory:

The current generation of computers provides us with sufficient memory per node and the aggregate memory for the codes MARS15, PHITS, MCNPX/6, COSY and LISE++. ANSYS is more stringent in that respect. The current memory requirement for ANSYS is 8 GB per node and 0.128 TB of the aggregate memory. It is expected that we will need 64 GB of memory per node and 1.024 TB of the aggregate memory by 2020, and 128 GB of memory per node and 2.048 TB of the aggregate memory by 2025.

3.4 Scratch Data and I/O:

MARS15, PHITS, MCNPX/6, COSY and LISE++ will not need a significant amount of scratch data or I/O bandwidth. The scratch space for ANSYS is 0.05 TB per calculation today, 0.4 TB per calculation in 2020, and 0.8 TB per calculation in 2025. No significant I/O bandwidth for ANSYS is needed.

3.5 Long-term and Shared Online Data:

The aggregate permanent online data storage for all the subprojects is 1 TB today, 10 TB in 2020, and 20 TB in 2025.

3.6 Archival Data Storage:

The aggregate archival data storage for all the subprojects is 1 TB today, 10 TB in 2020, and 20 TB in 2025.

3.7 Workflows:

Not addressed here.

3.8 Many-Core and/or GPU Readiness:

Our radiation transport and beam codes use one core per process. One process uses the entire model. We do not currently have code versions that are GPU-capable.

3.9 Software Applications, Libraries, and Tools:

PHITS, MCNPX/6, and MCNP data libraries are export-controlled codes. ANSYS is a commercial code.

3.10 HPC Services:

Not anticipated.

3.11 Additional Needs:

Not anticipated.

Requirements Summary Worksheet

Codes: MARS15, PHITS, MCNP6, COSY, LISE++, ANSYS	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	520,000	x6	x12
Computational node hours (homogeneous many-core)**	20,000	x8	x16
Computational node hours (with GPU or accelerator)	N/A	N/A	N/A
Memory per node	8 GB	64 GB	128 GB
Aggregate memory	0.128 TB	1.024 TB	2.048 TB
Data read and written per run	0.05 TB	0.4 TB	0.8 TB
Maximum I/O bandwidth needed	low	low	low
Percentage of runtime for I/O	low	low	Low
Scratch file system space needed	0.05 TB	0.4 TB	0.8 TB
Permanent online data storage	1 TB	10 TB	20 TB
Archival data storage needed	1 TB	10 TB	20 TB

* Core hours for conventional processors.

** Node hours for homogenous many-core architectures.

GRETA Streaming

Lead Author(s): M. Cromaz (mcromaz@lbl.gov) (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context:

GRETA (gamma-ray energy tracking array) is a large-scale, gamma-ray spectrometer to be sited at the Facility for Rare Isotope Beams (FRIB). GRETA's scientific reach extends among almost all FRIB scientific programs including nuclear astrophysics, nuclear structure, and fundamental symmetries. Much of the sensitivity of this array is derived from the extraordinary energy resolution these detectors provide. However, for a class of experiments that involve very high event rates, pileup limits the throughput of these devices unless compromises are made in shaping time and therefore resolution. However, having access to the full trace streamed to HPC resources (rather than now, an FPGA), one could apply considerably more sophisticated adaptive filtering algorithms. The I/O rates required for such a system are exceptional, and the development of such a system would require close collaboration with the high-performance networking and HPC communities.

1.2 Research Objectives for the Next Decade:

The research objectives are shared with those of GRETINA/GRETA. GRETA streaming, however, would enable a class of experiments that require very high gamma-ray counting rates at high throughputs. Due to the only recent availability of computer and networking hardware to facilitate such a device, we are in an early prototyping stage and would only expect such a system to be only deployed with the GRETA spectrometer that would fall into the 2025 time frame.

2. Computational and Data Strategies

2.1 Approach:

In the full streaming model the output of the GRETA digitizers (4000, 100 MHz, 16-bit) are directly output via standard networking to computing resources. Filtering is performed on these independent trace streams to generate energies, times, and trigger primitives. Trigger primitives are then forwarded to a global trigger-processing system where they are evaluated. Time stamps for valid events are then returned, and windowed traces are forwarded to the standard GRETA online processing system.

2.2 Codes and Algorithms:

Each trace stream will need to be forwarded to a core where multiple filters can be run. These filters could range from (simple) adaptive filters that alter their integration time based on the local environment to (complex) fitters that account for the detailed electronics response. Given the possibility of complex filters, it will be necessary to break a single trace amongst multiple cores. This could require redundant (overlapping) trace information to be sent as certain types of filters (such a baseline corrections) will require long histories. A mechanism to reassemble and filter out redundant information will be required.

3. Current and Future HPC Needs

3.1 Computational Hours:

Current - prototype - small-scale electronic prototyping

2025

$100,000 \text{ gamma} * 1 \text{ ms/core} = 100 \text{ cores/ch}$

$4,000 \text{ ch} * 100 \text{ core} = 400,000 \text{ cores}$

$3,000 \text{ hr} * 400,000 \text{ cores} = 1.2 * 10^9 \text{ core hours/yr}$

3.2 Parallelism:

The problem is naturally parallel on a channel level but a higher level of partitioning will be required if detailed fitting of the pulse trains are to be carried out. This will involve replication of data so that history effects can be taken into account. Coupling between tasks in a given channel involve communication of trigger primitives and identification of redundant results. We are not in the position to do this now, and it would require a program of targeted R&D. We will need high levels of consultation with groups experienced in the implementation of such systems.

3.3 Memory:

Memory requirements are determined by the amount of trace history for a given channel that is required and a scaling factor to account for storage of intermediate results. While the detailed algorithm has not yet been determined 0.2s of trace and a scaling factor of 3 are reasonable. Assuming that:

2025

$100 \text{ MHz} * 2 \text{ bytes/ch} * 0.2\text{s} * 3 = 120 \text{ MB/ch [node]}$

$4,000 \text{ ch} * 120 \text{ MB} = 480 \text{ GB [aggregate]}$

3.4 Scratch Data and I/O:

The I/O rate is determined by the streaming rate of digitizers. We believe that a factor of two in compression is achievable and can be implemented efficiently. Scratch data is not required by the streamer, or the GRETA spectrometer as a whole, if trace data is retained during the run.

2025 - max I/O rate

$100 \text{ MHz} * 2 \text{ bytes/ch} * 0.5 = 100 \text{ MB/s}$

$100 \text{ MB/s} * 4,000 \text{ ch} = 400 \text{ GB/s}$

3.5 Long-term and Shared Online Data:

In the context of a streamer, long-term and shared I/O in this case would be similar to that of the standard implementation of the GRETA spectrometer as only energies and interaction points are stored long term.

2025

10.3 PB

3.6 Archival Data Storage:

This is given by the GRETA spectrometer specification.

2025: 2.6 PB/yr

3.7 Workflows:

Like GRETINA/GRETA, GRETA streaming is a real-time system and the workflow is one of a data acquisition system. Data is passed from 120 digitizers via 40GB ethernet to computational resources. Trigger primitives are quickly extracted (~1s latency) and forwarded to trigger processors that identify valid events in the same stream. Simultaneously, high-resolution energies and times associated with each stream are extracted. These energies, times, and appropriate trace segments based on input from the trigger processor are then forwarded to signal decomposition. At that point the workflow is identical to that described for standard GRETA.

3.8 Many-Core and/or GPU Readiness:

This system is still in early development. It will be designed assuming an architecture involving processors with large numbers of cores or GPUs.

3.9 Software Applications, Libraries, and Tools:

This has not yet been determined. It is likely custom network drivers will be required.

3.10 HPC Services:

As mentioned in the introduction, the I/O rates required for a full streaming implementation are exceptional and would require close collaboration with the high performance networking and HPC communities.

3.11 Additional Needs:

As with GRETINA/GRETA, the requirements of reliability, availability, and stability over long time scales are essential to successfully carry out the experimental campaigns at FRIB.

Requirements Summary Worksheet

Code: GRETA streaming shares many values with GRETINA/GRETA tracking (do not “add” sheets directly)	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	NA	NA	1.2×10^9 core hr/yr
Computational node hours (homogeneous many-core)	NA	NA	NA
Computational node hours (with GPU or accelerator)	NA	NA	NA
Memory per node	NA	NA	0.12 GB
Aggregate memory	NA	NA	0.48 TB
Data read and written per run	NA	NA	3300 TB
Maximum I/O bandwidth needed	NA	NA	400 GB/sec
Percentage of runtime for I/O			
Scratch file system space needed	NA	NA	0 TB
Permanent online data storage	NA	NA	1030 TB
Archival data storage needed	NA	NA	2600 TB/year

* Core hours for conventional processors.

Gamma-ray Tracking for the GRETINA/GRETA Spectrometer

Lead Author(s): M. Cromaz (mcromaz@lbl.gov) (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context:

GRETA (gamma-ray energy tracking array) is a large-scale, gamma-ray spectrometer to be sited at the Facility for Rare Isotope Beams (FRIB). A smaller scale demonstrator of this array, GRETINA, is currently operating at the National Superconducting Cyclotron at Michigan State University (NSCL/MSU) and will move to ATLAS/ANL in 2017. This array is the first implementation of a gamma-ray tracking detector, which by employing complex signal-processing algorithms in real time, can locate the scattering path of incident gamma rays. This capability enables GRETINA/GRETA to perform high-resolution gamma-ray spectroscopy at very high efficiencies that gives it a scientific reach that extends among almost all FRIB scientific programs including nuclear astrophysics, nuclear structure, and fundamental symmetries.

1.2 Research Objectives for the Next Decade:

GRETINA currently operates a program that involves a wide range of nuclear structure and nuclear astrophysics experiments. Through 2021 the spectrometer will be used to carry out experimental campaigns at NSCL coupled to the S800 spectrograph and at ATLAS/ANL. During this time we will scale the number of detector modules and the counting rate capability. The GRETA spectrometer, a full 4 pi coverage implementation of GRETINA, is scheduled to be complete when FRIB begins operation in 2022. Its performance characteristics and computational needs are significantly higher than those of GRETINA.

2. Computational and Data Strategies

2.1 Approach:

GRETINA/GRETA achieves its very high sensitivity by being able to track the Compton scattering paths of multiple incident gamma rays from the reaction of interest. Key to this is the ability to resolve the location and relative energy depositions of gamma-ray interaction points within the Ge detector volume with high spatial accuracy. Computationally, this begins with the simulation of charge transport and signal generation within each segmented detector. A set of simulated signals, termed a basis, is calculated for unit charge on an equi-sensitivity grid that spans each detector volume. During the experiment, observed signals are then fit against linear combinations of these simulated basis signals to locate interaction points and their energies within the detector volume. Interaction points are then spatially grouped/clustered to assign them to likely gamma-ray hit and permutations of these points a fit to the Compton scattering formula to determine the most likely scattering sequence.

2.2 Codes and Algorithms:

The primary codes/algorithms for GRETINA/GRETA involve the determination of the location of interaction points. This task, known as signal decomposition, operates in two phases. The first phase is an adaptive grid search carried out on an equi-sensitivity grid that assumes two interaction points for each hit segment. Multiple hit segments are fit iteratively. This result is then used to seed a nonlinear fitting routine that allows for a third interaction, interpolation between grid points and refinement of

the time offset. These codes are run in a cluster that processes experimental data in real time (seconds level latency) to enable immediate analysis and refinement of experiments.

3. Current and Future HPC Needs

3.1 Computational Hours:

The computational hours required are given by the product of the experimental data rate and the running time of the instrument.

Current: 8 quads

10 ms/decomp for std core; 32 crystals @ 1 Khz = 32k decomp/s, which gives 320 core * s

1500 hours / year - $320 * 1500 = 4.8 \times 10^5$ core hours / year

2020 - high rate, 12 quads

12 ms/decomp (ext. basis); 48 crystals @ 1.5 kHz = 72k decomp/s, which gives 864 core * s

1500 hours / year - $864 * 1500 = 1.3 \times 10^6$ core hours / year

2025 - GRETA at FRIB

100 ms/decomp (rough estimate); 120 crystals @ 4 kHz = 480k decomp/s, which gives 48,000 core * s

3000 hours/year - $48,000 * 3000 = 1.4 \times 10^8$ core hours/year

3.2 Parallelism:

The code uses course-grained parallelism (process level). In the future we plan to modify the parallelization model to admit limited communication between decomposition processes to relay global state information of the event and to efficiently utilize emerging, large-scale multi-core processors.

3.3 Memory:

Memory requirements are dominated by the simulated signal basis set.

Current: $1.6 \text{ GB} * 2 = 3.2 \text{ GB} / \text{node}$; 32 crystals * 3.2 GB = 102 GB aggregate

2020: $1.6 * 1.2 * 2 = 3.8 \text{ GB} / \text{node}$; 48 crystals * 3.8 GB = 182 GB aggregate

2025: Very uncertain; depends on memory bandwidth, need for local redundant copies; Given a model of 100 cores/node with a single basis image: $1.6 (1.2) \text{ GB} = 1.9 \text{ GB/core}$, $500 * 1.9$ indicates 960 GB aggregate

3.4 Scratch Data and I/O:

In GRETINA scratch data currently refers to raw trace data that is written to disk for diagnostic purposes. Interaction points are employed in experimental analysis and written to disk. In GRETA, although this is not yet defined, it is possible that raw trace data will be written to disk as well to allow for offline signal decomposition to refine interaction point positions. We assume an experimental run to be of 10-day duration. Below are figures for both input I/O (trace) where the network I/O rate is maximum, scratch storage for a given run, and output I/O (interaction point, trace in 2025).

Maximum I/O rates (trace):

Current: $16\text{kB/tr event} * 32\text{k/s decomp} = 0.5 \text{ GB/s}$

2020: $16\text{kB/tr event} * 72\text{k/s decomp} = 1.2 \text{ GB/s}$

2025: $8 \text{ kB/tr event} * 480 \text{ k/s decomp} = 3.8 \text{ GB/s}$

Scratch:

Current: $0.5 \text{ GB/s} * 3600 * 240 \text{ hr/run} = 432 \text{ TB}$

2020: $1.2 \text{ GB/s} * 3600 * 240 \text{ hr/run} = 1036 \text{ TB}$

2025: 0 TB – as trace data is written no scratch space is req'd

Stored:

Current: $0.5 \text{ kB/event} * 32 \text{ k/s decomp} = 16 \text{ MB/s}; 16 \text{ MB/s} * 3600 * 240 \text{ hr/run} = 13.8 \text{ TB / run}$

2020: $0.5 \text{ kB/event} * 72 \text{ k/s decomp} = 36 \text{ MB/s}; 36 \text{ MB/s} * 3600 * 240 \text{ hr/run} = 31.1 \text{ TB / run}$

2025 – GRETA: $8 \text{ kB/ tr event} * 480 \text{ k/s decomp} = 3.8 \text{ GB/s}; 3.8 \text{ GB/s} * 3600 * 240 \text{ hr/yr} = 3.3 \text{ PB / run}$

3.5 Long-term and Shared Online Data:

In the context of GRETINA/GRETA, long-term I/O is defined by the length of time data needs to be accessible until it is analyzed and/or replicated for analysis by the experimental group responsible for the data. Assuming an accessibility window for the data of 3 months and running an average of 1500 hr/year in the current and 2020 timeframes, and 3000 hrs/yr in the 2025 timeframe the storage requirements are as follows:

Current: 21.6 TB

2020: 48.6 TB

2025: 10.3 PB

3.6 Archival Data Storage:

For GRETINA/GRETA archival data storage for a certain class experimental data is required by DOE and NSF mandates. Currently all mode 2 (interaction point data) falls under this mandate. In the current and 2020 time frames, this number is the same as the long-term storage rate scaled for the year.

Current: 86.4 TB/yr

2020: 195 TB/yr

2025: $0.5 \text{ kB/evt} * 480 \text{ k/s decomp} = 240 \text{ MB/s}; 240 \text{ MB / s} * 3600 * 3000 \text{ hr / yr} = 2.6 \text{ PB/yr}$

3.7 Workflows:

GRETINA/GRETA process their experimental data in (near) real-time [$< 10\text{s}$]. Raw, windowed trace data from ~1000 [GRETINA] / ~4000 [GRETA] digitizer channels is forwarded to a set of nodes that perform a signal-conditioning step followed by signal decomposition to determine interaction points within the HPGe crystal volume. Interaction points from each crystal are then sent to a global event builder where, based on a time-stamping mechanism, they are assembled into global events along with information forwarded by auxiliary detector systems. Once global events are formed, they can be forwarded to nodes that run the Compton tracker, which groups and orders interaction points in gamma-ray hits. Summary statistics are then formed and presented to experimenters online to evaluate and optimize the experiment.

3.8 Many-Core and/or GPU Readiness:

Currently the software for GRETINA is not ready for a transition to large numbers of lightweight cores or hardware accelerators. We will need to transition to these platforms in order to scale to GRETA rates at acceptable cost levels. This will be addressed in the context of the GRETA project. We will need high levels of consultation with groups experienced in the implementation of such systems.

3.9 Software Applications, Libraries, and Tools:

Currently the libraries and tools used by the Gretina project are standard and well supported among Linux distributions. The only exception to this is the use of EPICs-based state machines in the computing cluster to manage the online signal-processing tasks so the control system can be unified. For GRETA (2025) we expect not to have such a dependency and move to a more standard distributed control model.

3.10 HPC Services:

None.

3.11 Additional Needs:

GRETINA/GRETA require a very high degree of reliability, availability, and stability over long time scales to successfully carry out the experimental campaigns at FRIB.

Requirements Summary Worksheet

Code: GRETINA/GRETA Signal Tracking	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	4.8 x 10 ⁵ / yr	X 2.7	X 290
Computational node hours (homogeneous many-core)	NA	NA	NA
Computational node hours (with GPU or accelerator)	NA	NA	NA
Memory per node	3.2 GB	3.8 GB	1.9 GB
Aggregate memory	0.10 TB	0.18 TB	0.96 TB
Data read and written per run	13.8 TB	31.1 TB	3300 TB
Maximum I/O bandwidth needed	0.5 GB/sec	1.2 GB/sec	3.8 GB/sec
Percentage of runtime for I/O			
Scratch file system space needed	430 TB	1,030 TB	3,280 TB
Permanent online data storage	21.6 TB	48.6 TB	1,030 TB
Archival data storage needed	86.4 TB/yr	195 TB/yr	2600 TB/yr

* Core hours for conventional processors.

FRIB Spectrometers

Lead Author(s): D. Bazin, R. Zegers (zegers@nscl.msu.edu) (National Superconducting Cyclotron Laboratory, Michigan State University) and A. Couture and S. Mosby (Los Alamos National Laboratory)

1. Description of Research

1.1 Overview and Context:

Nuclear physics experiments that involve rare isotope beams at various incident energies use different reactions to probe the evolution of the nuclear force across the nuclear landscape. The particular areas covered by this case study involve the use of magnetic spectrometers such as the S800, HRS (High Rigidity Spectrometer), and ISLA (Isochronous Spectrometer with Large Acceptances) to filter, collect, and analyze reaction products. These experiments involve cutting-edge electronics, data acquisition systems, and analysis/simulation software that heavily rely on high-end computing, storage, and networking technologies.

1.2 Research Objectives for the Next Decade:

The scientific goals for the next decade in this domain of research cover a wide variety of themes that warrant the need for a recoil spectrometer to track and identify reaction residues. These themes range from in-beam gamma-ray spectroscopy to invariant-mass spectroscopy of unbound states in exotic nuclei. The scientific program is prioritized by the nuclear science community and guided by the recommendations in the 2015 NSAC Long Range Plan.

The computational and data analysis goals can be divided in two categories:

- Online data taking and processing: multiple data acquisition systems working in parallel need to be monitored and synchronized, as well as collect data. The requirements for this category will likely fall in the realms of network infrastructure, I/O performance and capacity for real-time processing of high-volume data.
- Offline data analysis: the processing of the data will require resources in the domains of CPU and I/O performance. Parallel processing could offer increased performance but only if performed at a high level (parallel processing of individual events). Similarly, gains in performance are sought for simulation purposes, where parallel processing is important to simulate large experiments with sufficient statistics.

2. Computational and Data Strategies

2.1 Approach:

The data acquisition rates remain the main limitation, attributed to the performance of the electronics. In the next decade, the expectation is that electronics will be updated to modern technologies. This will, in turn, decrease the dead time and at the same time increase the data throughput. As higher data-taking rates are available and the electronics evolves more toward digitization of the signals, the burden

on the CPU and network needs will increase. Consequently, additional capacity for real-time processing and storage will increase significantly. Parallel processing of online experiments is envisioned.

2.2 Codes and Algorithms:

The current codes and work flows are dedicated to event merging between different parts of the data acquisition, as well as online data processing for online monitoring and visualization of the experiments. These tasks are presently accomplished by the NSCL-supported DAQ package (see contribution from R.L. Varner). This package will evolve and must be able to meet the online processing requirements of future experiments. This includes the increased demand for support of DAQ systems with digital electronics, which comes with higher data throughput (10–100 times greater than experiments based on analog DAQ systems).

3. Current and Future HPC Needs

3.1 Computational Hours:

Not addressed here.

3.2 Parallelism:

Not addressed here.

3.3 Memory:

Not addressed here.

3.4 Scratch Data and I/O:

Not addressed here.

3.5 Long-term and Shared Online Data:

Not addressed here.

3.6 Archival Data Storage:

Not addressed here.

3.7 Workflows:

Not addressed here.

3.8 Many-Core and/or GPU Readiness:

Not addressed here.

3.9 Software Applications, Libraries, and Tools:

Not addressed here.

3.10 HPC Services:

Not addressed here.

3.11 Additional Needs:

Not addressed here.

Requirements Summary Worksheet

I am trying to estimate numbers in this table for one experiment. It is envisioned that the HRS, ISLA, and S800 spectrometers will be used for at least 50% of all experiments at FRIB. We assumed that one experiment is 100 hrs.

Code: FRIB Spectrometers	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	100	X10	X20
Computational node hours (homogeneous many-core)**	100	x10	X20
Computational node hours (with GPU or accelerator)***	0	†	†
Memory per node	16 GB	64 GB	256 GB
Aggregate memory	0.05 TB	1 TB	2T B
Data read and written per run (= experiment)	0.5 TB	5T B	10 TB
Maximum I/O bandwidth needed	0.1 GB/sec	1 GB/sec	5 GB/sec
Percentage of runtime for I/O	50%	50%	50%
Scratch file system space needed	0.5 TB	5T B	10 TB
Permanent online data storage	1 TB	10 TB	20 TB
Archival data storage needed (cumulative for all experiments)	50 TB	500 TB	1,000 TB

* Core hours for conventional processors.

** Node hours for homogenous many-core architectures.

***Node hours for GPU or accelerator usage.

†Not yet decided; GPUs might be useful for trace analysis of digital data.

CUORE

Lead Author(s): Tommy O'Donnell (Virginia Polytechnic Institute and State University)

1. Description of Research

1.1 Overview and Context:

The Cryogenic Underground Observatory of Rare Events (CUORE) is an international collaboration that aims primarily to search for neutrinoless double-beta ($0\nu\beta\beta$) decay. This is a lepton-number-violating process predicted to occur by some extensions of the Standard Model of Particle Physics that have been developed to accommodate non-zero neutrino masses. Discovery of $0\nu\beta\beta$ decay would prove that lepton number is not conserved by nature and establish that neutrinos are Majorana fermions. Moreover, it would support theories that leptons seeded the matter–antimatter asymmetry in the Universe. The potential for fundamental impact motivates intense effort to search for this decay.

The CUORE experiment is a large array ($\sim 1,000$) of bolometers operating essentially as individual detectors at close to 10 mK. High-performance computing (HPC), storage, and networking are essential and timely tools enabling the scientific program of CUORE. In particular, access to adequate computing power is needed for real-time monitoring of the health and stability of the detectors, to thoroughly analyze the collected data, and to perform detailed simulations necessary to constrain systematic uncertainties associated with our apparatus and methods. Furthermore, as we proceed to physical interpretation of our data, evaluating the statistical significance of our conclusions can be a computationally intensive undertaking, especially with such rare-event data where common approximations are inadequate. Access to high-density and fast-access storage is essential for data management and data stewardship, while high-bandwidth networks are vital for sharing data across our multinational collaboration.

1.2 Research Objectives for the Next Decade:

In the next decade we aim to fully complete the science program of CUORE and exploit the close to 4 ton-years of TeO_2 exposure we will accumulate. Our goals include a search for—or perhaps discovery of— $0\nu\beta\beta$ decay of ^{130}Te , a precision measurement of the $2\nu\beta\beta$ decay half-life and spectral shape, searches far beyond the Standard Model processes such as axion interactions, and searches for dark matter. Furthermore, we will exploit the data to measure and identify the background sources in the detector materials and surrounding environment to unprecedented precision. This later goal will require extensive simulations but will be vital to identifying the most promising design and materials selection strategies to further improve the physics reach of this type of detector.

Anticipating deployment of new ancillary high-bandwidth cryogenic light detectors that are currently in development, we expect our data footprint and computational needs to increase 5-to 10-fold, but exploiting this technology will reduce our largest background by approximately 100-fold.

2. Computational and Data Strategies

2.1 Approach:

CUORE is an array of approximately 1000 detectors that function largely individually. Some challenges are listed below.

1. Effective real-time monitoring of the full array. We currently meet this challenge using a web-based interface that allows detector operation experts to dynamically query diagnostics variables from a MongoDB database and display the results. Some development of these dynamic display tools is outsourced to specialized companies, while the diagnostic quantities are determined by detector experts.
2. The data footprint of CUORE is approximately 100 TB/year raw data. This is relatively modest compared to some other nuclear physics experiments. We plan in addition to maintain a full data backup on NERSC's HPSS archive system. This is also modest relative to existing capacity.
3. Final analysis and physics interpretation of the energy spectrum from the CUORE detectors involves high multi-dimensional fit to the data. This fit is repeated a large number of times on simulated data sets to understand the statistical significance of our conclusions and estimate systematic uncertainties. We rely on parallelization to accomplish this analysis in a reasonable time.
4. Detailed Monte Carlo simulations that reproduce particle trajectories and energy deposition patterns in our detectors are vital to modeling backgrounds. Again we solve this problem with parallelization.

2.2 Codes and Algorithms:

The work flow involves the following steps: raw data collection, software event triggering, rapid approximate online analysis, population of MongoDB with diagnostic variables, raw data copy to local HPC cluster at experiment host facility (LNGS), full event reconstruction analysis (data quality selection, signal filtering, pulse height evaluation, pulse-height to energy calibration, data blinding, coincidence tagging), high-level event selection and spectral analysis for $0\nu\beta\beta$ decay search, and other studies. Population of postgres database with operations data and summary analysis variables. Distribution of raw data to remote backup sites and distribution of high-level data summary files to collaborating institutions. Our codes are largely written in C++ and build off of the ROOT framework. We rely on python and perl for scripting and task automation.

Simulation tools, event generators, and particle tracking are based on the Geant4 toolkit. A phonon simulation in Geant4 is in the early stages of development.

3. Current and Future HPC Needs

3.1 Computational Hours:

See table.

3.2 Parallelism:

We rely on course-grain parallelism (up to 300x). We do not have current plans for fine-grained parallelism.

3.3 Memory:**3.4 Scratch Data and I/O:****3.5 Long-term and Shared Online Data:****3.6 Archival Data Storage:****3.7 Workflows:****3.8 Many-Core and/or GPU Readiness:**

No. We are interested in exploring these possibilities. Training would be desirable.

3.9 Software Applications, Libraries, and Tools:**3.10 HPC Services:**

- Advanced training
- Desirable if workable skills-set can be acquired in the typical training cycle of advanced graduate students/postdocs
- Security
- Visualization assistance
- Gateways

3.11 Additional Needs:

Requirements Summary Worksheet

Code: CUORE	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	106	2x	5x
Computational node hours (homogeneous many-core)			
Computational node hours (with GPU or accelerator)	0	0	0
Memory per node	2–8 GB	1x	2x
Aggregate memory	TB	TB	TB
Data read and written per run	0.3 TB	1.5x	3x
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O	20	1x	1x
Scratch file system space needed	100 TB	2x	4x
Permanent online data storage	500 TB	2x	5x
Archival data storage needed	2000 TB	2x	10x

* Core hours for conventional processors.

The ALICE Experiment at the LHC¹⁰

Lead Author(s): R. Jefferson Porter (Lawrence Berkeley National Laboratory)

1. Description of Research

1.1 Overview and Context:

The physics research conducted by the ALICE experiment focuses on the study of the quark-gluon plasma (QGP), a phase of strongly interacting matter formed at the extreme energy densities in heavy ion collisions at energies reachable at the Large Hadron Collider (LHC). The experiment takes large data samples from PbPb, pPb, and pp collisions with the ALICE detector and relies on distributed computing for data processing, simulations, and data storage, organized within the Worldwide LHC Computing Grid (WLCG)¹¹ collaboration.

1.2 Research Objectives for the Next Decade:

The experiment is in its second three-year data-taking cycle, providing first ever measurements at the highest energies achieved in heavy-ion collisions. The research expands on that done at lower energies, detecting identified particles at low to moderate transverse momenta (pT) and extends the investigation to higher pT regions attainable at LHC energies. ALICE operates approximately 75,000 CPU cores along with an approximately 50 PB disk-resident data store, in which both CPU and disk storage are distributed at its 80 grid sites around the world.

The next phase of the experiment, starting in 2020, will focus on measuring low-pT heavy-flavor and charmonium production for which interactions with the QGP medium are calculable. The physics program leverages the large heavy-flavor production rates at LHC energies, which can provide large statistics within minimum-bias event samples. The experiments will run the detector in a continuous (triggerless) readout mode and compress the data within each event in real time, so as to maximize the event sample without loss of physics information. A project referred to as O2 (Online-Offline)¹² is under way in ALICE to develop a new, more flexible framework that will allow ALICE to meet its physics objectives at extremely large data rates.

2. Computational and Data Strategies

2.1 Approach:

The major computational and data challenge for ALICE is managing its large-scale distributed data processing. This challenge is met today by a workload management system, AliEn¹³ (Alice Environment), developed by the ALICE and CERN IT, built on central services for global job and data management and lightweight site services that interact with the central services to efficiently map processing tasks to available compute and data resources. The ALICE workload is distributed, largely independent of job type, onto its set of nearly homogenous facilities.

¹⁰ <http://alice-collaboration.web.cern.ch/>

¹¹ <http://wlcg.web.cern.ch/>

¹² <https://cds.cern.ch/collection/ALICE%20Reports?ln=en>

¹³ <http://alien.web.cern.ch/>

2.2 Codes and Algorithms:

ALICE software uses a ROOT-based framework¹⁴ for I/O and data processing. The codes span a broad range of pattern recognition algorithms used to extract physical signals from background and interpret physical characteristics from those signals. Codes are written by physicists familiar with the science objectives, augmented by ROOT libraries for data handling and common analytic techniques in statistical analysis and visualization. In addition to data processing, ALICE relies on a series of codes to simulate physics signals and evaluate detector responses using the GEANT¹⁵ tool set to simulate the behavior of collision products as they interact with the detectors. Over half of ALICE computing resources is used by simulation tasks.

3. Current and Future HPC Needs

3.1 Computational Hours:

ALICE computing requirements are organized in the WLCG-accepted benchmark of HEPspec06 (HS06), for which a conventional core purchased in 2014 is approximately 13 HS06. The minimal CPU requirement for ALICE-distributed processing is about 600,000 HS06, or about 45,000 conventional CPU cores, providing about 400 million CPU hours during the year. The real-time data processing in online is run on more heterogeneous architectures and accounts for another 50–60 million CPU hours each year. These requirements will approximately double over the next four years. After that, design changes in the experiment will make very large data rates available in about 2020 and will require a dramatic change in the ALICE computing model.

The future computing model will rely on a much larger processing capacity in online used for real-time data reduction and first-level processing. A dedicated facility is being designed for this purpose to be deployed at the experiment site. In addition, simulation and data analysis will continue to be distributed to remote resources, though some of those resources may be targeted for specific tasks. The best estimates are that the processing needs outside of the online facility are an order of magnitude larger than today's needs, or about 5 billion CPU hours per year.

3.2 Parallelism:

ALICE event-based processing is embarrassingly parallel with each event treated as distinct and small enough relative to the memory per core to be processed independently. Fine-grained parallelism is limited to only the real-time data processing done online, in which GPU or other techniques make use of the parallel structures from detector segmentation. The detector segmentation should allow fine-grained parallelism in future simulation work, once those capabilities exist in underlying GEANT software used by ALICE and other HEP/Nuclear Physics experiments.

Coarse-grained parallelism is not currently being used in ALICE but expected to play a larger role in the O2 framework under development. A critical feature of the framework is the capability to allow different processing tasks at the sub-event level to operate independently, with I/O managed by message services and separate tasks dedicated to distributing work and collating results.

¹⁴ <https://root.cern.ch/>

¹⁵ <http://www.geant4.org/geant4/>

3.3 Memory:

The current ALICE computing model has a series of requirements placed on the processing at the level of a single job core: 2 GB of memory, 10 G of local disk space, and external network connection. Since many tasks are built on a common ALICE-specific code base, processes often have significantly larger total memory footprints than needed while running. As a result, there is an additional recommendation to allow another 2–4 GB of swap space per core. ALICE software does not make use of either shared memory or any memory hierarchy; however, this should change with the future O2 framework.

3.4 Scratch Data and I/O:

The data and I/O needs are best described separately for simulations and data analysis. Simulations are CPU-intensive, where typical jobs run for several hours, require little input, and generate modest sized GB-scale output files used in later analysis jobs. Both scratch disk volumes and I/O requirements are estimated on a per-core basis of about 10GB and 1MB/s. In contrast, data analysis is generally I/O-intensive, using input data sets of about 50TB that are processed repeatedly. They need to be kept on long-term, online storage. The analysis codes need to ingest input data at about 5MB/s per core in order to keep the fraction of time in I/O down to a reasonable amount.

3.5 Long-term and Shared Online Data:

The primary long-term online data storage is needed for analysis-ready reduced data, known as Analysis Object Data (AOD). While current AOD data sets are on the order of 50 TB, there are many variants of individual data sets. Thus a useful computing facility on the order of 1,000 cores requires PB-scale long-term online storage. The requirement will grow significantly in the 2020 time frame, during which analysis will likely be targeted to specific analysis facilities, possibly one located in the US. Such an analysis facility will need to have at least 5 PB of online disk capacity in the 2020 time frame, rising to 15–20 PB by 2025.

3.6 Archival Data Storage:

The ALICE experiment has extensive requirements for archival storage on Tier-1 facilities; however, the current plan for ALICE-US does not include providing a Tier-1 facility.

3.7 Workflows:

ALICE operates a distributed grid facility of approximately 80 grid sites, the majority of which are Tier-2 facilities on which both simulations and data analysis jobs are run. Each site consists of a compute cluster, one or more grid-enabled storage elements, and a set of site-level services run on a single node, the ALICE VObox. The ALICE workload management system is built on central services, a global File Catalog, and task queue from which work is pulled to the site in a work flow known as a “pilot-job” model. That is, the site services submit lightweight “JobAgents” onto the local cluster, which contacts the central services to request a payload job that runs as a forked process. The workload management system sends jobs that require input data, such as analysis jobs, to the sites where the input data exists on the local storage elements. Output data is stored both on the local storage element and distributed onto additional storage elements via an AliEn algorithm.

In the O2 era of 2020, the extremely large data samples will be more efficiently analyzed on dedicated analysis facilities with large online storage elements (10s PBs) to hold all AOD data sets and large CPU capacity for processing demands of ongoing analyses. ALICE expects to have a small number of such facilities available, perhaps with one in the US.

3.8 Many-Core and/or GPU Readiness:

The future ALICE framework being developed as part of the O2 project is required to be more modular and flexible relative to the ALICE ROOT-based framework used today. Such a framework may allow better use of future lightweight cores. Special cases where the software is optimized for specific architectures, such as GPUs, are in well-defined online pattern recognition routines. We expect community-wide R&D in common ROOT-based algorithms and GEANT-like simulation codes will be the most efficient way to develop new routines that take advantage of these new technologies.

3.9 Software Applications, Libraries, and Tools:

ALICE software relies on the ROOT analysis framework for data I/O and processing and on the GEANT tool set for simulations. In the future, the O2 framework will also incorporate a messaging service, such as OMQ, as well as new database technologies for managing its experimental conditions data, now handled in a ROOT-based file system.

Software deployment is managed through the CERN VM File System (CVMFS), which consists of a network of repository stores fed by a central source repository, from which the experiment software is accessed on each worker node from a dynamic cache mounted as a Fuse-based read-only file system. The CVMFS system requires each worker node to have external network, local disk space, and Fuse installed in the system.

ALICE uses XRootD¹⁶ as the current disk-based storage technology for managing its distributed data sets on the ALICE grid. The grid infrastructure is currently being migrating to use the CERN-developed EOS¹⁷ system, which adds a high-performance data management structure on top of an XRootD back end. ALICE expects that its long-term distributed data store solution will be based on the evolution of the CERN EOS system.

All ALICE grid monitoring of job and data operations, in conjunction with compute, storage, and network resources, is through the MonALISA¹⁸ infrastructure. All sites and services, from the central servers to the individual job processes, are monitored with MonALISA.

Access for ALICE collaborators to its large CPU and storage facilities is enabled in a secure way via a direct interface to user-level job submission and data management through the AliEn system. User access is provided through X.509 certificates registered in the ALICE VOMS system. ALICE-USA participants may obtain certificates from the Open Science Grid (OSG) through the ALICE VO registration agent.

ALICE Grid resources are provided to ALICE as part of agreements with the WLCG collaboration. The US-based OSG is a peer grid to the WLCG and provides the WLCG with accounting reports on site utilization. The ALICE-USA computing facilities operate OSG gateway services that process jobs and provide accounting of those jobs back to the WLCG. In addition to the above software tools, ALICE uses MySQL for its central services, Git for code management, Jenkins for software build services, and Jira for bug tracking.

¹⁶ <http://xrootd.org/>

¹⁷ <https://eos.web.cern.ch/>

¹⁸ <http://monalisa.caltech.edu/monalisa.htm>

3.10 HPC Services:

The ALICE work flow places several site-level requirements that are historically nonstandard on HPC systems. These include continuous operation of edge services (e.g., ALICE VObox), external (outgoing) network connection to the WAN, and local disk space for scratch and swap, and software delivery with CVMFS. The distributed ALICE work flow does require federated identity services that are recognized at every grid site. Special arrangements have been made for ALICE at its U.S. sites, LLNL, LBNL, and ORNL to support its workload management. It would benefit the project if such agreements were based on DOE-sanctioned identity federations acceptable to all sites.

Requirements Summary Worksheet

The following requirements are filtered by the expected US contribution to ALICE, which is estimated to be less than 10% of ALICE totals. Along with this order-of-magnitude reduction, several components of the work flow such as online data reduction, event reconstruction, and archival storage, are not expected to be US contributions to ALICE computing operations and are omitted here.

ALICE Processing	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	80 million	5x	10x
Computational node hours (homogeneous many-core)	0	30 million	50 million
Computational node hours (with GPU or accelerator)	0	10 million	20 million
Memory per node	150 GB	300 GB	300 GB
Aggregate memory	300 TB	3000 TB	6000 TB
Data read and written per run	100 TB	4000 TB	8000 TB
Maximum I/O bandwidth needed	5 GB/sec	100 GB/sec	200 GB/sec
Percentage of runtime for I/O	30%	30%	30%
Scratch file system space needed	70 TB	500 TB	1,000 TB
Permanent online data storage	1,000 TB	5,000 TB	15,000 TB
Archival data storage needed	0 TB	0 TB	0 TB

Hall-D/GlueX at Jefferson Lab

Mark Ito and David Lawrence (Hall D, Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

The GlueX Experiment in Hall D was built primarily to search for exotic mesons predicted by Quantum Chromodynamics (QCD), the theory of the strong interaction. In particular, the search focuses on states where exotic quantum numbers demonstrate gluonic degrees of freedom, those beyond those exhibited by well-known quark–anti-quark mesons. Such states are known as hybrid mesons. Evidence for these states has been reported previously, but the results have not been unambiguously confirmed.

The experiment exploits major components of Jefferson Lab's recently completed 12-GeV upgrade: (a) increase of the electron beam energy from 6 GeV to 12 GeV and (b) a new experimental hall for conducting experiments with a tagged coherent photon beam, producing a high flux of 9-GeV photons

Since GlueX was first proposed, advances in lattice QCD calculations have predicted a whole spectrum of hybrids, with both exotic and conventional quantum numbers. To firmly establish the existence of hybrids will require observation of multiple states of this spectrum.

1.2 Research Objectives for the Next Decade:

To establish a spectrum of hybrid mesons and explore other related aspects of hadron spectroscopy using these new tools will require the greater part of the next decade. There are also other physics topics that can be addressed with the data to be taken. All will require the computational resources outlined below.

2. Computational and Data Strategies

2.1 Approach:

To carry out the program, many billions of particle collisions will have to be recorded at a data rate of roughly 300 MB/s. Individual events are about 20 kB. The computational problems are of several types.

1. **Calibration.** The responses of each of the detector elements must undergo a calibration process based on statistical analysis of detector response measured against known physical quantities. This involves multiple parameters of tens of thousands of electronic channels.
2. **Offline Reconstruction.** Recorded events will go through a reconstruction stage where low-level data from individual detector elements are combined to form a global picture of the collision. Reconstructed events give high-level properties of detected particles. Reconstructed events are recorded in an appropriate format, much more compact than the original “raw” data.
3. **Offline Analysis.** The reconstructed data is analyzed repeatedly and in different ways to obtain physics results. Results are almost exclusively obtained from various types of statistical analyses of the data.

4. **Simulation.** Since the observation of collisions is not generally complete and always of finite resolution, the effect of the detector on recorded events must be assessed. The most powerful tool for this is to simulate collisions in terms of their individual detector responses and perform the same reconstruction and analysis on the simulated events as is done with the real data. This must be done on a sample of events comparable to the real data to give reliable results.

2.2 Codes and Algorithms:

We use several independent software packages:

- sim-recon: programs and libraries for performing event reconstruction and physics analysis
- [JANA](#): framework for multi-threaded processing of event-based data
- HDDS: Detector geometry specification
- [HDDM](#): The Hall D Data Model, a format for event-based data.
- [EVIO](#): The CEBAF Online Data Acquisition ([CODA](#)) event format.
- [CCDB](#): Calibration Database
- [RCDB](#): Run Conditions Database
- [AmpTools](#): Amplitude analysis package (aka partial wave analysis)
- [EventStore](#): a package for managing and deploying data files and event lists.

3. Current and Future HPC Needs

3.1 Computational Hours:

Our plans call for about 90 million core hours per year. We do not have projections for how this might grow, but a factor of 2 every five years is a conservative estimate.

3.2 Parallelism:

For reconstruction, which is the major consumer of CPU time, we run multi-threaded on single nodes, typically 24 threads. Jobs can be run in parallel on independent nodes; during production we can run a few thousand jobs at a time.

3.3 Memory:

Our jobs require several GB of memory to run; 5 GB is a typical number, 10 GB in some cases.

3.4 Scratch Data and I/O:

We estimate 200 TB of scratch space, mainly for staging data in and out of a tape library, both raw and reconstructed.

3.5 Long-term and Shared Online Data:

We estimate about 300 TB of reconstructed data and post-reconstruction data, always available from disk, between 2020 and 2025.

3.6 Archival Data Storage:

So far we have stored 500 TB on our tape library. We expect that grow at a rate of roughly 1 PB per year for the foreseeable future.

3.7 Workflows:

There are two distinct workflows:

1. **Batch Processing.** Using computing at Jefferson Lab and off-site processing farms a series of jobs are submitted, each job differing only in the event data being processed. This includes calibration, reconstruction, and large-scale simulation.
2. **Interactive/Near-Interactive processing.** Using computing resources local to the user, analysis of reconstructed data or post-reconstruction data to do physics analysis. Processing is done repeated with variations in algorithms on a particular set of event data.

3.8 Many-Core and/or GPU Readiness:

Event-based data allows parallelization as described in Section 3.2, but runs into practical limitations on GPU-like platforms. Independent events can be processed independently, but the algorithm will branch differently from event to event. Sub-event parallelization is possible, but is not something that is commonly done in the field; it presents a new set of challenges.

One exception for us is the AmpTools package. It takes advantage of a particular features of the technique used to fit collections of events with quantum mechanical amplitudes with free parameters that lend itself to a GPU-like architecture. The AmpTools packages can run user-supplied CUDA-based code to run on standard NVIDIA GPUs.

3.9 Software Applications, Libraries, and Tools:

- CLHEP [<http://proj-clhep.web.cern.ch/proj-clhep>]
- Doxygen [<http://www.stack.nl/~dimitri/doxygen/>]
- Geant4 [<http://geant4.cern.ch/>]
- Git [<https://git-scm.com/> Git]
- MariaDB [<http://www.mariadb.com/>]
- ROOT [<http://root.cern.ch/>]
- Scons [<http://www.scons.org/>]
- SQLite [<http://www.sqlite.org/>]
- Xerces [<http://xerces.apache.org/> Xerces].

3.10 HPC Services:

No needs for non-standard HPC services are anticipated.

3.11 Additional Needs:

Robust WAN connection for distribution of data among collaborators at institutions outside JLab will be important to facilitate collaboration and distribution of human effort and use of remote computing facilities.

Requirements Summary Worksheet

Code:	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	90 million	x2	x4
Computational node hours (homogeneous many-core)	N/A	N/A	N/A
Computational node hours (with GPU or accelerator)	N/A	N/A	N/A
Memory per node	5–10 GB	Same GB	Same GB
Aggregate memory	5–10 TB	TB	TB
Data read and written per run	1,000 TB	Same TB	Same TB
Maximum I/O bandwidth needed	1 GB/sec	Same GB/sec	Same GB/sec
Percentage of runtime for I/O	1%	Same	Same
Scratch file system space needed	200 TB	Same TB	Same TB
Permanent online data storage	50 TB	500 TB	500 TB
Archival data storage needed	1,000 TB	5,000 TB	10,000 TB

* Core hours for conventional processors.

CLAS12 Physics Data Processing

Lead Author(s): H. Avakian, G. Gavalian, V. Gyurjyan, and V. Ziegler (Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

The detailed understanding of the orbital structure of partonic distributions, encoded in Transverse-Momentum Dependent (TMD) and generalized parton distributions (GPDs), has been widely recognized as one of the key objectives of the JLab 12 GeV upgrade project and driving force behind construction of the Electron Ion Collider. A model-independent extraction of TMDs and GPDs from semi-inclusive and hard-exclusive DIS data is the critical component of future data analysis. The cross-sections of different semi-inclusive and exclusive processes will be extracted in multi-dimensional space, requiring new approaches for extraction, storage, and analysis of that data.

1.2 Research Objectives for the Next Decade:

The main goal of future studies of the nucleon are precision measurements of spin and azimuthal asymmetries in semi-inclusive pion and kaon and exclusive photon production off unpolarized, longitudinally polarized and transversely polarized targets, providing access to the spin and flavor dependences of transverse-momentum and space distributions of quarks in the valence region. Development of a 3D PDFs extraction framework with validation procedures will be crucial for accomplishment of that program. In particular, multi-dimensional storage of cross-sections in elementary bins limited in size by detector resolution will require very significant increases in MC simulations of multi-dimensional detector acceptance for processes relevant for studies of the 3D structure of the nucleon. The focal new features of CLAS12 include operation with a luminosity of 10^{35} $\text{cm}^{-2} \text{sec}^{-1}$ and improved acceptance and particle detection capabilities at forward angles. In order to keep up with the expected data rates at this luminosity, we need data-acquisition and data-processing systems capable of processing data in quasi-real time. This approach assumes data in-memory stream processing that will undoubtedly require faster networks, faster and bigger memories, and more computing resources to satisfy our data-taking and data-processing needs for the next decade.

2. Computational and Data Strategies

2.1 Approach:

For the CLAS12 data processing—where data volumes are predicted to increase substantially, simulated data is geographically distributed across collaborating Universities, information changes swiftly, and expertise is widely dispersed—we need a different approach to CLAS12 physics data processing and data provisioning. The future of the physics-computing ecosystem is predicted to move from being computer-centric to more data-centric. We have to change our approach and use architectures that are designed to process data in near real time. That might mean doing basic analyses of data while it is in the memory or doing data processing in a network as the data is flowing (stream processing).

2.2 Codes and Algorithms:

CLAS6 data-processing applications were mostly single-threaded and extremely slow. In order to increase data-processing speeds, we deployed multiple applications in a batch-processing environment, making entire systems extremely inefficient and I/O limited. What we really need for CLAS12 is a customizable and elastic (multi-threaded and horizontally scalable) component-based software system that can perform data stream processing, minimizing I/O usage and data migrations. Data-processing algorithms and code solutions must be packed in autonomous components (application building blocks) that can communicate among each other through data-message passing. Each component can be developed, tested, and optimized independently by different collaborating groups, making CLAS12 software collaboration and software contribution more effective. This component-based environment will give much needed agility and efficiency to CLAS12 overall data processing.

3. Current and Future HPC Needs

3.1 Computational Hours:

The current estimate of the number of core days for reconstruction, validation studies, analysis, calibration, and simulations is estimated to be 11,605 core days per running day to keep up with the expected event rate under normal running conditions. About 77% of this number comes from Geant 4 simulations assuming a ratio of number of simulated events to triggered data event of 10:1. By 2020 we expect to have plans in place to increase the luminosity possibly by a factor of 2. Data acquisition rates are expected to increase multi-fold due to a new readout-blocking capability of the CODA: JLAB data acquisition system.

3.2 Parallelism:

Currently the CLAS12 data-processing framework (CLARA) linearly scales vertically (utilizing multi-core architecture CPUs) and horizontally (JLAB batch farm and AWS). The CLARA framework implements service-oriented architecture (SOA) in a flow-based programming (FBP) paradigm to enhance the efficiency, agility, and productivity of data-processing applications. CLARA is an approach to developing data-processing applications based on the concept of multiple asynchronous processes (software building blocks), called services, communicating by means of data streams. Thus, an application is viewed as a system of data streams being transformed by services.

3.3 Memory:

Currently our data processing has 2 parallel running work flows: data processing and data provisioning. The data-provisioning work flow is responsible for staging data files into the memory (RAM disk) and removing results of the memory. Our present requirement is to have an on-board memory big enough to be able to stage multiple data files in in-memory data pipe. At this time, the latency to process a single event is larger than memory access latency. However, instead of presently used event-level parallelism, we are planning to use sub-event level parallelism in the future, which will require faster memory access.

3.4 Scratch Data and I/O:

We are planning to parallelize disk I/O and decouple it from the data-processing work flow. We are planning to perform in-memory data stream processing as much as is reasonably possible. Estimated online scratch storage space will be ~100 TB by 2020 and ~500 TB by 2025.

3.5 Long-term and Shared Online Data:

We would like to keep an active sample of our data for online access., Based on the CLAS6 data-mining project in which one run can take 2 TB data, CLAS12-produced data will require 4 TB data per run period, 8 TB in 2020 and 50 TB in 2025.

3.6 Archival Data Storage:

3.7 Workflows:

CLAS12 data-processing work flow consists of data, and service layers, distributed across multiple multi-core computing nodes. The data layer itself is divided into sublayers, presented by different types of data, such as experimental, simulated, or analyzed (e.g., data summary tape) data. The service layer houses a library of data-processing modules/services that can be used to process data. Complex data-processing applications can be designed by algorithmically combining multiple services together. We expect services to function close to the data source, prevented data migration. The future data processing application design and execution, as well as data access and provisioning we plan to perform through the CLARA front end, making user data geographic location agnostic. The system's front end will house a meta-data database that describes not only stored data but also services/processes used to generate the data. Front-end also accommodates the meta-data database for library processes/services, thus playing the role of the application server. For data visualization and monitoring, CLARA web tools will be utilized.

3.8 Many-Core and/or GPU Readiness:

The flow-based programming paradigm is well suited for the lightweight multi-core systems. This is one of the main drivers to use CLARA as a base framework for our data-processing needs. However, as the number of cores increases, to deal with the power envelope, cores become less complex, providing decreased single-thread performance. This will definitely force us to design CLARA services and make them responsible only for processing a small fraction/part of the event (sub-event level parallelization). We expect to invest effort to redesign some of the currently used data-processing algorithms. This is not an easy task since fragmentation of some algorithms (e.g., kalman filter) is not obvious.

3.9 Software Applications, Libraries, and Tools:

The CLARA service-layer layer sits on the top of the *service-bus* that provides an abstraction of the locally developed publish-subscribe messaging system (xMsg). xMsg is using a ZeroMQ socket library and implements messaging patterns like topic pub-sub, workload distribution, and request-response. So ZeroMQ is one of the third-party software packages that we are using. The third-party packages used are all open source. In the future we might need more efficient IPC (interprocess communication) packages. Note that the CLARA architecture allows integration of the new technologies without painful rewrite or redesign of the entire system. This approach provides an application designer the ability to modify data-processing applications by incorporating different services in order to find optimal data-processing conditions, thus demonstrating the overall agility of the CLARA framework.

To aid detector reconstruction code development, we have created a unified common tools library. The tools include I/O libraries for streaming data through CLARA services and dictionary-based bank deserialization. A common geometry package was developed to implement all detectors in a unified framework and provide 2D and 3D detector visualization to detector reconstruction code developers. A newly developed data visualization package is used for analyzing the data and developing online monitoring and calibration tools.

3.10 HPC Services:

Data accessibility and data intelligence will increase predictive power of the data and transform our minds and revolutionize our understanding of nature in the 21st century. For data accessibility we need tools to uniformly describe and categorize data at least within the nuclear research communities. The tools to insure data security will be increasingly important in the transparent data environment.

Requirements Summary Worksheet

Code: _____	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)*	2.5x10 ⁸	x2	x2
Computational node hours (homogeneous many-core)			
Computational node hours (with GPU or accelerator)			
Memory per node	100 GB	100 GB	100 GB
Aggregate memory	300 TB	600 TB	600 TB
Data read and written per run	TB	TB	TB
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O			
Scratch file system space needed	TB	TB	TB
Permanent online data storage	TB	TB	TB
Archival data storage needed	TB	TB	TB

* Core hours for conventional processors.

D.3 Case Studies Addressing Nuclear Structure and Reactions

Nuclei from First Principles

Lead Authors: G. Jansen (Oak Ridge National Laboratory), J. Vary (Iowa State University), and S. Gandolfi (Los Alamos National Laboratory)

Contributing Authors: Hai Ah Nam (Los Alamos National Laboratory), E. Ng (Lawrence Berkeley National Laboratory), and P. Maris (Iowa State University)

1. Description of Research

1.1 Overview and Context:

We compute properties of atomic nuclei consisting of nucleons (i.e., protons and neutrons) by solving the quantum many-body problem using inter-nucleon interactions as the only input. The interactions are tuned to reproduce selected scattering and bound-state properties, and the ensuing *ab initio* or first-principles calculations yield predictions for properties of rare and stable isotopes that can be confronted with experiments. As recognized by previous DOE reports, exascale computations are required to fulfill the critical role of making predictions for and guiding new experiments, and important advances are required that solve challenging computational issues arising with changes in architectures.

1.2 Research Objectives for the Next Decade:

The Facility for Rare Isotopes (FRIB), under construction at Michigan State University, will study the physics of neutron-rich nuclei that are important for the creation of all the heavy elements in the universe. To address the needs of FRIB, we must be able to reach computationally the extremes of stability for neutron-rich nuclei, study their structure and decay properties, and provide quantified uncertainties on our predictions. Equally important are the plans to search for neutrinoless double-beta decay; see the case study “Electroweak Phenomena.” To help formulate a successful search strategy and to interpret the data, we need to extend *ab initio* solutions of nuclei to medium and heavy nuclei using nucleon–nucleon (NN) plus three-nucleon (3N) interactions. To accomplish this, we must be able to work with larger distributed data sets, achieve more parallelism from our work flow and calculations, and develop higher order algorithms for better accuracy and precision. This will require leadership-class production runs to accommodate the storage needs and the inter-node communication loads. In addition, we require improvements to the sparse matrix multiplication algorithms to achieve improved computational efficiency through increased parallelism.

2. Computational and Data Strategies

2.1 Approach:

We employ configuration interaction (CI), coupled cluster (CC), and the Auxiliary Field Diffusion Monte Carlo (AFDMC) methods to solve the quantum many-body problem numerically. The CI approach sets up a sparse matrix eigenvalue problem for the nuclear Hamiltonian in a chosen basis representation and produces the low-lying eigenvalues (nuclear spectra) and the eigen functions (nuclear wave functions). With nuclear strong interactions and increasing nucleon number, the method requires more memory

than available on current machines to go above about mass $A = 20$ and obtain solutions reasonably converged with respect to many-body basis truncation. We need to reach much larger nucleon numbers to address the research objectives. Fortunately, new theoretical methods where the *ab initio* Hamiltonian is re-expressed as an equivalent Hamiltonian acting only among the valence nucleons and opening the path to heavier systems using a sequence of CI calculations that are achievable on exascale architectures. The CC method currently addresses nuclei near-closed shells and closed subshells. It scales polynomially in the number of nucleons and makes controlled approximations to the quantum mechanical Schrödinger equation. The main computational tasks are contractions of spherical tensors that are large and sparse data structures. The tensors are distributed in memory guided by load balance, and parts of the tensors are duplicated to minimize the need for communication and synchronization. The quest for increased precision and for heavier nuclei requires us to modify our CC computational strategy, because it will no longer be possible to duplicate enough data to avoid significant communication. The AFDMC algorithm solves the Schrödinger equation with Markov chain Monte Carlo to sample the path integral used to project out the ground state. These branching random walks employ significant linear algebra at each step of the simulation. It scales polynomially in the number of nucleons, and it produces exact results for a given local Hamiltonian. Extensive statistics are needed for large systems because of the sign problem. In addition, nontrivial load balancing to control the distribution of configurations among nodes is needed. Most of the computing time is spent diagonalizing matrices with dimensions in the few hundreds, and this has to be repeated billions of times to obtain good statistics.

2.2 Codes and Algorithms:

The nuclear coupled-cluster code (NUCCOR) is a nuclear physics tool running on Titan at OLCF today. It takes a set of nuclear two- and three-body interactions as input and solves the many-body Schrödinger equation approximately by constructing similarity transformations based on the solutions of nonlinear systems of equations, and subsequently diagonalizes the similarity transformed Hamiltonian in a truncated basis. The similarity-transformed Hamiltonian is in general non-Hermitian, large, sparse, and structured, and a small subset of extremal and interior eigen pairs is needed. The calculations are performed with multiple different interactions to properly extrapolate to the final results. Convergence decisions are made manually, and additional calculations are scheduled as needed. The MFDn (Many-Fermion Dynamics–nuclear) configuration interaction (CI) code has played a leading role in pushing the boundaries of the atomic masses accessible through No-Core Shell Model (NCSM) calculations, including 3N interactions. MFDn solves the nuclear many-body problem through diagonalization of a large, sparse Hamiltonian matrix. It evaluates and stores the sparse symmetric many-body Hamiltonian matrix in lower triangular form over all available nodes using a compressed row-column format. We employ a Lanczos iteration scheme with reorthonormalization of vectors at each iteration. A low-lying set of eigenvalues and eigenvectors is obtained and used to evaluate observables for comparison with experiments. We perform a sequence of calculations with MFDn beginning in smaller basis spaces (smaller matrix dimension) and proceed up to the largest feasible basis space and then extrapolate to the infinite matrix limit for the final result. We repeat this process to map the dependence on the basis space parameters and measure uncertainties. In addition, variations in the extrapolations provide additional uncertainties in the prediction. A typical case solves for the lowest 15 eigenvalues and eigenvectors with matrix dimensions ranging from 1 billion to 25 billion. Significant benefits have been derived from improvements in the sparse Hamiltonian construction and matrix-vector multiplication schemes within the locally optimal block preconditioned conjugate gradient within the eigen solver. The AFDMC code has been used to study medium-mass nuclei and nuclear matter with NN and 3N interactions. It uses branching random walks to simulate the strongly interacting many-body system and

scales polynomially rather than exponentially with the number of nucleons because of the auxiliary fields used to sample the spins and isospins.

3. Current and Future HPC Needs

3.1 Computational Hours:

For a single run, NUCCOR (MFDn) scales well up to 100,000 (300,000) conventional cores on Titan at OLCF. MFDn scales well up to 500,000 cores on MIRA. NUCCOR and MFDn each use 20–40 million core hours on Titan yearly, and MFDn uses 25–25 million core hours on MIRA yearly, all through an INCITE allocation. We expect our requirement for each application to increase by 10X by 2020 and by another 10X by 2025. Going to a higher precision for NUCCOR and heavier nuclei for MFDn requires 3–4 orders of magnitude additional floating point operations. We expect 1–2 orders of magnitude better performance from a combination of algorithmic improvements and advances in the many-body theory. For a single run, the AFDMC code scales well up to about 130,000 cores on MIRA at ANL, and similarly at NERSC, and yearly uses 10–20 million core hours aggregated between Institutional Computing provided by LANL and NERSC. We expect our requirement to increase by a factor of 20 by 2020 and by another 10–50 by 2025, depending upon available resources.

3.2 Parallelism:

NUCCOR exploits multiple layers of parallelism in solving the nuclear many-body problem. At the coarsest level, many almost-independent calculations with modest dependencies are run in parallel. A team of MPI ranks performs threaded calculations on a distributed data set. The most compute-intensive kernels are run on GPU accelerators, and the list of accelerated kernels is increasing as part of the CAAR project at OLCF. We plan to exploit additional parallelism by integrating more of our work flow as coarse-grained tasks, with the possibility of launching additional tasks as part of a decision-making process. In addition, we plan to create computational kernels optimized for different architectures that can be used interchangeably and in parallel for heterogeneous systems. We have developed code segments for MFDn that efficiently utilize the GPUs on Titan for the time-consuming process of evaluating the Hamiltonian matrix with $3N$ interactions. We have implemented a compressed block storage scheme that improves OpenMP scalability at larger numbers of OpenMP threads/MPI process, which performs well on MIRA. With the NESAP award at NERSC, we are achieving additional improvements in MFDn by improving overlapping communications and computation and by increased parallelism within each node. Currently, AFDMC simply uses MPI tasks; we have implemented OpenMP as well. The latter is likely to play a more significant role as we move forward to very many-core architectures. It is likely we will also have to change load balancing to ADLB or a similar library that manages quasi-independent tasks.

3.3 Memory:

Currently NUCCOR requires a minimum of 8 GB of shared memory on a node for data duplication as part of a communication-avoiding strategy. We are limited by the aggregate memory available on the full system, which we suspect will be the case also the next 10 years. We are currently exploring the use of multi-level memory hierarchies where the slower levels are used as read-only cache to store non-zero matrix elements. In addition, we are exploring the possibility of generating matrix-elements on-the-fly, exchanging aggregate memory-requirements for additional floating-point operations. MFDn distributes the many-body Hamiltonian matrix among all the nodes used in the run without duplication. Memory inefficiencies and limits arise due to the need to store the input $NN + 3N$ matrix elements on each node as well as copies of the Lanczos vectors. This need arises due to the communication costs that would

otherwise explode with increasing node counts. To achieve our research goals, we will benefit from increasing memory on each node as this helps reduce overall memory inefficiency (i.e., reduces duplicated information). AFDMC tremendously simplifies the memory requirements for each instance of the random walk. The memory needs are likely to grow significantly, however, as we incorporate more sophisticated algorithms to deal with the fermion sign problem and the dynamic response. This will require us to explore the use of the high-bandwidth memory and burst buffer in new generation machines.

3.4 Scratch Data and I/O:

Currently, NUCCOR writes at most 15 TB of checkpoint/restart data to scratch space for a single run and reads at most 50 TB of input data, while producing at most 2 TB of result data. Using the current checkpoint/restart strategy, we will write 1,500 TB of checkpoint/restart data during a run (because of next-order excitation operators) and produce 200 TB of result data due to using a higher order excitation operator requiring 2 orders of magnitude more storage, while reading 1,000 TB of input data. We can tolerate 5% of total time runtime spent doing I/O, more if I/O operations are nonblocking and can be performed by a background thread. Assuming a runtime of 24 hours, the required bandwidth will increase to 600 GB/sec. Algorithmic improvements and more reuse of data will reduce this requirement. We will also investigate the possibility of storing compressed data sets to reduce the total storage requirement and bandwidth. Most runs of MFDn use input/output files on the order of 1–50 GB. Large production runs produce 1 TB wave function output files for post-processing; large runs using 3N interactions use 10- GB input files. We use MPI I/O and concurrent I/O-computation to mitigate associated I/O costs. Regular partial checkpointing occurs for the largest runs, requiring a few TBs of scratch space. The AFDMC can easily do checkpoints by writing the configurations at a given imaginary time on the disk. The output for ground-state calculations is usually limited to few hundred MB for a given ensemble of configurations. At this point, the initial and final I/O for a run requires only a few percent of the total run time. This will likely grow significantly with advanced algorithms and response and may require several TB for storage and may require the use of parallel I/O.

3.5 Long-term and Shared Online Data:

Today, NUCCOR (MFDn) uses around 100 TB (50 GB) of input data for 3N interactions, which will increase to 2,500 TB (1 TB) in 2025, mainly due to the need for higher fidelity calculations and uncertainty quantification (UQ). In addition, we plan to share the development of matrix element data files for effective operators associated with a wide class of observables. These additional matrix element files have the potential to exceed the sizes of the interaction files by 1–2 orders of magnitude due to their non-scalar character. All these matrix element files could be used by the entire nuclear physics community and should be a shared resource. Nuclear physics calculations are performed on all major computing facilities in the US, which requires high-bandwidth connections between the facilities. With AFDMC to large systems, where the propagation in imaginary time is expensive, we plan to store configurations in such a way that other observables can be post-processed, as commonly done in lattice QCD calculations. Certain data should be stored long term and made available online for later analysis.

3.6 Archival Data Storage:

NUCCOR and MFDn each currently produce up to 2 TB of data for a production run, which largely consists of wave function components. These are archived not only for verification purposes but also for later use in calculating observables we cannot compute today. This will increase to 200 TB by 2025.

3.7 Workflows:

QMC (GFMC and AFMC) work flows are fairly simple, with large-scale path-integral calculations checkpointed occasionally. Analysis is often run simultaneously with the calculation of the path integrals. Alternatively, storage of intermediate configurations is possible for calculations of additional observables and history. In IMSRG, separate binaries are chained together by scripts. In the first stage of GCM, HFB basis states are generated. These are stored and read in as input for computation of matrix elements.

3.8 Many-Core and/or GPU Readiness:

In **NUCCOR**, we exploit hardware accelerators in the form of GPUs for the most compute-intensive kernels by creating abstractions that can have implementations optimized for different architectures. We are currently working on a nonblocking execution with overlapping computation and data transfer as much as possible. Our plan is to extend this to a fully task-based execution model, where the task scheduler is locality aware and can make sure data is in the correct memory before a task is executed. We think this model will work well for both lightweight cores and hardware accelerators with deep memory hierarchies. We have developed code segments for MFDn that efficiently utilize the GPUs on OLCF's Titan for the time-consuming process of evaluating the Hamiltonian matrix with $3N$ forces. We are working on developing additional tasks for the GPUs that currently reside on the CPUs by, for example, implementing the locally optimal block preconditioned conjugate gradient method within the eigen solver. At present, we have implemented MPI with OpenMP for the AFDMC code. The memory and communication requirements are limited so we can run efficiently using one MPI process per core. We expect this to change for many-core hardware. The code is already using a task-based execution model, which, together with the modest memory requirements, should allow us to run effectively on the many-core architectures. To the extent possible, we will use local tasks to minimize required movement of data

3.9 Software Applications, Libraries, and Tools:

We rely on highly optimized BLAS and LAPACK libraries as well as the parallel HDF5 library for I/O. Also, we require compilers that support the current FORTRAN 2003/2008 standards as well as the upcoming FORTRAN 2015 standard. We use and plan to use various standard tools for version control, debugging, profiling, building, regression and unit testing, and bug and feature tracking. We will require tools for debugging and profiling that are efficient and practical to use at exascale. Abstractions and tools for performance portability will become increasingly more important as computer architectures evolve. While no definitive needs have been identified at this stage, visualization of the results may become important.

3.10 HPC Services:

We will have an increasing need for smaller scale systems to perform pre- and post-processing.

3.11 Additional Needs:

We estimate a need for additional 3 domain/computational scientist positions over 5 years in this area. These new positions are critical to enable the effort to scale to the largest scale machines heading toward exascale.

Requirements Summary Worksheet

Code: MFDn, NUCCOR, AFDMC	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	150 million	0.5X	0.5X
Computational node hours (homogeneous many-core or GPU)	0	400 M	3,000 M
Memory per node	36 GB	128 GB	128 GB
Aggregate memory	750 TB	3.7 PB	9 PB
Data read and written per run	100 TB	400 TB	2 PB
Maximum I/O bandwidth needed	15 GB/sec	50 GB/sec	500 GB/sec
Percentage of runtime for I/O	5	5	5
Scratch file system space needed	20 TB	1,000 TB	5,000 TB
Permanent online data storage	100 TB	1,000 TB	5,000 TB
Archival data storage needed	20 TB	1,500 TB	5,000 TB

Note: future requirements are listed in node hours (not core hours) for GPU and multi-core.

Few-Nucleon Reactions

Lead Authors: S. Quaglioni (Lawrence Livermore National Laboratory),
R. Schiavilla (Old Dominion University and Thomas Jefferson National
Accelerator Facility)

Contributing Authors: Petr Navrátil (Tri-University Meson Facility,
University of British Columbia)

1. Description of Research

1.1 Overview and Context:

One of the outstanding challenges in contemporary nuclear theory is to describe nuclear structure, clustering, and reactions within a single coherent framework, in which the nucleus' constituents, the nucleons, interact via realistic nuclear forces and electroweak currents. Besides its intrinsic intellectual merit relating to our understanding of the strong and electroweak interactions in nuclei, such a coherent framework is required in order to enable accurate predictions for difficult-to-measure reactions that synthesize the elements in stars or come into play in terrestrial applications of nuclear fusion. It is also critical for improving our understanding of the structure of weakly bound nuclei, and for using nuclei as probes of fundamental physics through, for example, beta decay, double-beta decay, and neutrino-nucleus scattering. Bolstered by high-end computing, advanced quantum Monte Carlo (QMC) techniques and new efficient approaches, such as the ab initio no-core shell model with continuum (NCSMC), have made great strides towards meeting this challenge in light nuclei (up to ^{12}C) by enabling the first ab initio calculations of a variety of strong and electroweak processes within the realistic framework. Further developments and exascale resources will be critical to improve the accuracy and extend the range of applicability of these methods to heavier systems and more complex reactions.

1.2 Research Objectives for the Next Decade:

Our target problem for the next decade will be to arrive at a quantitative description, within NCSMC, of binary and ternary reaction processes in systems with up to $A = 20$ nucleons, among which the $2\alpha(\alpha,\gamma)^{12}\text{C}$ and $^{12}\text{C}(\alpha,\gamma)^{16}\text{O}$ reactions critical for the formation of life and the synthesis of heavier elements in stars.

2. Computational and Data Strategies

2.1 Approach:

In the NCSMC, we solve the quantum-mechanical problem of the dynamical interactions of two (three) nuclear clusters made of strongly interacting neutrons and protons. We first solve the microscopic Hamiltonian eigenvalue problem for each of the fragments and the aggregate system by working in a many-body harmonic oscillator basis. We then combine the static solutions for the aggregate system with continuous "microscopic-cluster" states, made of pairs (or triplets) of nuclei in relative motion with respect to each other, to recover the full dynamics of the system. The relative wave function between (among) nuclear fragments and the scattering matrix are obtained by solving a set of nonlocal coupled-channel integral-differential equations. In order to achieve the research objectives, we will need to greatly reduce the present cost of computing the interactions of nuclear targets with nucleons and deuteron projectiles starting from chiral nucleon–nucleon (NN) plus three-nucleon (3N) forces, and

extend such capability to further describe collisions involving $3\text{H}/3\text{He}$, as well as interactions among three nuclear fragments. This will require a critical improvement in the scalability of our codes, as well as advanced techniques for distributed storage in memory and quick access of stored quantities, such as 3N-force matrix elements.

2.2 Codes and Algorithms:

NCSMC calculations are carried out using a suite of HPC codes, using MPI and OpenMP parallelism, by the name FUSION (Fundamental Unified Structure and Interactions of Nuclei). Its major components/work flow can be described as follows: MANYEFF+V3TRANS (and its two-body version, NCSMPN2BEFF) computes effective matrix elements of the nuclear Hamiltonian used to find for few- and many-body solutions as well as few-body translational invariant wave functions of the projectiles (angular momentum and linear algebra); NCSD (or BIGSTICK, MFDn) computes Slater-determinant (SD) target and aggregate-nucleus many-body wave functions (bit manipulations, hashing algorithms, importance truncation, Lanczos matrix diagonalization); TRDENS computes one- and two-body densities of the wave functions and/or parts of the SD couplings required (bit manipulations, hashing algorithms, angular momentum algebra); NCSM_RGM uses one- and two-body densities and/or SD interaction couplings, as well as projectile wave functions, to compute translational invariant projectile-target interactions and coupling form factors with the aggregate nucleus. It solves for projectile-target relative motion wave function and scattering matrix (angular momentum algebra, real and complex linear algebra).

3. Current and Future HPC Needs

3.1 Computational Hours:

Much of the NCSMC work so far has been performed on LLNL computers, supported by the Institutional Computational Grand Challenge project led by S. Quaglioni. In 2015 we used approximately 340M hours on VULCAN, a 5-petaFLOP/s Power7, PowerPC A2 (IBM) machine with 16 processors/node and 16-GB memory per node, for a total of 393,216 compute cores. We further used ~50M hours on Cray XK7 TITAN machine within the INCITE project led by Vary. In 2015, our usage increased by almost a factor of 8 compared to our more modest rate of about 1.2 in the preceding years. This was due to the extension of our calculations toward more complex reactions with heavier nuclear targets/projectiles, and the transition to VULCAN, where we need to use a larger number of nodes to reach the aggregate memory requirements of our calculations. We anticipate that in the next 10 years we will work to implement more efficient algorithms and memory distribution and to improve the parallelism of our codes to further optimize CPU usage. At the same time, given our scientific goals of roughly doubling the size of the systems with respect to what is currently accessible and perform calculations with heavier projectiles and for three-fragment reactions, we expect that our CPU usage will keep ramping up to reaching on the order of 1.5B and 4.0B hours in 2020 and 2025, respectively.

3.2 Parallelism:

NCSM_RGM uses a hybrid MPI/OpenMP parallelization as well as MPI/IO. As of today, it has been run with as many as 98,304 cores (composed of 6,144 nodes, with 12,288 MPI tasks with 8 OpenMP threads per MPI task). TRDENS is based on the MPI-2 protocol and was run with up to 8,192 cores on VULCAN. In the coming years we plan to work with computer scientists to radically improve the scalability of our codes and explore the use of accelerators and emerging technologies.

3.3 Memory:

We currently require a minimum of 16 GB of memory per node, but would benefit from the maximum feasible memory per node, which would allow us to increase the model-space size and hence the fidelity of our calculations. Our requirements for the aggregate memory grow with the size of the system under study.

3.4 Scratch Data and I/O:

Typically, a few hundred GB of scratch data per calculation is used, of which about half is restart data (our application has several built-in restart procedures). The other half (input, such as interaction matrix elements and wave functions, and output such as Hamiltonian and norm couplings, scattering phase shifts, scattering wave functions, as well as elements of the scattering matrix) can be moved to long-term storage and utilized for different applications and/or benchmark purposes and for post-processing of reaction observables. We expect that our data requirements will grow in the next decade, particularly in input. We expect that the output should not exceed 1 TB, while we estimate that we will reach on the order of 5 TB of input per run, mainly due to the increased size of the target and composite-nucleus eigenvectors. We use MPI I/O to mitigate costs associated with large input files. We would like to keep the time devoted to I/O to less than 5% of the total runtime.

3.5 Long-term and Shared Online Data:

Currently, we do not use online long-term storage, but it would be very useful in the future to share data (input interaction matrix elements, output scattering matrix elements, etc.) among collaborators in the US, Canada, and Europe.

3.6 Archival Data Storage:

About 2 TB of HPSS storage at the Livermore computing facilities is currently used. We expect that this will grow to about 50 TB by the year 2025.

3.7 Workflows:

The current work flow for FUSION is described in Sec. 2.2. In the next years, we plan to gradually phase out the use of intermediate densities in the computation of the projectile-target interactions in favor of precomputing SD interaction couplings. This will be essential for increasing the size of systems we can study.

3.8 Many-Core and/or GPU Readiness:

For the time being we have not explored the use of GPUs, mainly because for our computational problem, it is more straightforward to use OpenMP capabilities, and in part due to the lack of an adequate workforce and the lack of expertise to explore these new technologies. We envision that GPUs may be advantageous during the solution of the two- and three-body dynamical equations, where the calculation of the scattering matrix requires dense linear algebra (matrix multiplications and inversions, eigenvalue problems) at each energy step. Our strategy for exploiting GPUs and other new technologies will be to seek the help of applied math and CS experts.

3.9 Software Applications, Libraries, and Tools:

We need to link our codes to the multi-threaded Intel Math Kernel Library (MKL). Our compiler of preference is the Intel FORTRAN compiler. We make use of Lustre file system software to stripe large input files onto multiple hard drives and speed up MPI I/O. At the moment we have implemented our own binary file formatting through MPI I/O, but we may transition to HDF5 in the next years. Abstractions and tools for performance portability will become increasingly more important as

computer architectures evolve. While no definitive needs have been identified at this stage, visualization of the results may become important.

3.10 HPC Services:

We would benefit from consulting and account support, data analytics assistance, advanced training, and collaboration tools.

3.11 Additional Needs:

We estimate a need for additional 2 domain/computational scientist positions over 5 years in this area. These new positions are critical to enable the effort to scale to the largest-scale machines heading toward exascale.

Requirements Summary Worksheet

Code: FUSION	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	360M	0.5X	0.5X
Computational node hours (homogeneous many-core or GPU)		460M	2500M
Memory per node	8 GB	128 GB	128 GB
Aggregate memory	150 TB	3,000 TB	5,000 TB
Data read and written per run	1 TB	2 TB	6 TB
Maximum I/O bandwidth needed	15 GB/sec	0.05 GB/sec	0.1 GB/sec
Percentage of runtime for I/O	5%	5%	5%
Scratch file system space needed	1 TB	100 TB	400 TB
Permanent online data storage	Not used	1,000 TB	3,000 TB
Archival data storage needed	2 TB	1,000 TB	3,000 TB

Electroweak Phenomena

Lead Authors: J. Carlson (Los Alamos National Laboratory), J. Engel (University of North Carolina), R. Schiavilla (Old Dominion University and Thomas Jefferson National Accelerator Facility)

Contributing Authors: S. C. Pieper (Argonne National Laboratory)

1. Description of Research

1.1 Overview and Context:

The weak, electromagnetic, and strong interactions are inextricably linked and provide a unique tool to understand the underlying nuclear structure and dynamics. These tools can then be exploited to investigate neutrino physics and the role of weak interactions in astrophysics and fundamental symmetries. We will thus investigate electroweak phenomena in a large set of nuclei over a broad range of energies and momenta. At low energies, this research will provide accurate quantitative estimates of nuclear matrix elements that are critical for the interpretation of neutrinoless double-beta decay ($0\nu\beta\beta$) experiments, will further our knowledge of weak transitions relevant to the creation of the elements and tests of physics beyond the Standard Model, and will let us understand, quantitatively, neutrino propagation in the dense matter in core-collapse supernovae and neutron stars. At higher energies, it will provide essential input for unraveling the response of neutrino detectors in experiments, such as the Deep Underground Neutrino Experiment (DUNE), which will measure neutrino oscillation parameters, including the charge-parity (CP) violating phase in the neutrino mass mixing matrix that makes neutrinos and antineutrinos oscillate differently, and will determine the neutrino mass hierarchy.

1.2 Research Objectives for the Next Decade:

We plan to carry out *ab initio* calculations—solutions of the nuclear many-body problem with NN and 3N effective interactions between and among neutrons and protons and with associated weak currents—of (i) weak transitions between low-lying states of light and medium-weight nuclei, such as ${}^{16}\text{O}$, ${}^{40}\text{Ar}$, and ${}^{48}\text{Ca}$, in Quantum Monte Carlo (QMC), (ii) electroweak response functions describing inclusive low- (few to tens of MeV) and high-energy (hundreds of MeV to less than a GeV) neutrino scattering induced by neutral and charge-changing weak currents off nuclear targets, such as ${}^{12}\text{C}$ and ${}^{40}\text{Ar}$, and (iii) low-energy neutrino propagation in nucleonic matter, in QMC, and (iv) the matrix elements for $0\nu\beta\beta$ in ${}^{76}\text{Ge}$, ${}^{136}\text{Xe}$, and other nuclei that could be used in $0\nu\beta\beta$ experiments, in coupled cluster (CC) theory, the in-medium similarity renormalization group (IMSRG), and the generator-coordinate method (GCM).

2. Computational and Data Strategies

2.1 Approach:

For nuclei with mass number A up to 16, we plan to use Green's function Monte Carlo (GFMC) to project out of trial states the exact ground- and low-lying-states of these systems by application of the evolution operator in imaginary time. This method accounts explicitly for the complex spin and isospin dependence of nuclear potentials and currents, and is therefore limited to $A = 16$ by the exponential growth in computational requirements with the number of nucleons. For heavier nuclei, two other QMC methods are available: auxiliary-field diffusion Monte Carlo (AFDMC) and cluster variational Monte Carlo

(CVMC). In AFDMC the spin degrees of freedom are sampled along with the spatial ones in the stochastic evaluation of the imaginary-time propagation, while CVMC relies on a cluster expansion for the noncentral correlations in the evaluation of expectation values and transition matrix elements. The extension of these methods to the calculation of the response functions proceeds in two steps: first, the evaluation of their Laplace transforms (Euclidean response functions) with methods similar to those used in projecting out the ground- and low-lying states from trial states, and second, in the inversion of these Laplace transforms by maximum-entropy methods to obtain the actual response functions.

CC theory provides nonlinear algebraic equations for a similarity transformation of the Hamiltonian to a convenient non-Hermitian form. A small subset of interior eigen pairs of the large, sparse Hamiltonian matrix must then be obtained. The IMSRG integrates 108–109 coupled first-order ordinary differential equations (ODEs) for matrix elements of a unitary transformation operator that decouples low-lying eigenstates from the rest. The Magnus exponential expansion is used to represent the transformation and apply it to other operators. Transformed Hamiltonians are subsequently used as input for large-scale eigenvalue problems. The inclusion of continuum degrees of freedom and increasing emphasis of the multi-reference character will necessitate the transition to generalized Hermitian or complex symmetric eigenvalue problems for the treatment of excited states. The GCM is a variational method based on nuclear DFT that expands the full nuclear wave function in a set of nonorthogonal quasiparticle vacuums, each having fixed values for a set of collective coordinates (deformation, pairing, etc.). It can be used by itself or to provide a starting point for subsequent IMSRG evolution. We calculate not only spectra and $0\nu\beta\beta$ matrix elements but also two-body densities (and later three-body densities) for use in the IMSRG.

2.2 Codes and Algorithms:

For QMC (**GFMC** and **AFDMC**), the algorithms are similar to those used for calculations of nuclear structure. They involve branching random walks with large sparse (GFMC) or dense (AFDMC) linear algebra at each step of the walk. For **CC**, see Case Study “Nuclei From First Principles.” In **IMSRG**, the Magnus expansion allows the use of lower order ODE solvers (e.g., Euler or low-order Runge-Kutta). The eigenvalue problems that allow access to excited states currently employ Lanczos/Arnoldi methods. Matrix-vector products are expressed in terms of tensor contractions. In **GCM**, individual matrix elements evaluated through Gaussian quadrature and matrix determinants/Pfaffians.

3. Current and Future HPC Needs

3.1 Computational Hours:

QMC currently uses 200M–500M core hours on all applications. A typical response calculation for ^{12}C at one momentum transfer requires ~20M core hours on MIRA to yield sufficient accuracy to invert the response. Between 5 and 10 different momentum transfers are required for both neutral and charged-current response, requiring on the order of several hundred million core hours total. Further resources will be required for beta decay, for $A = 14$ in GFMC and for heavier isotopes using AFDMC. CC currently uses 300M–400M hours total on all applications. Double-beta decay in ^{76}Ge will require 50M hours, times 5 or 10 for uncertainty quantification. The time for ^{136}Xe would be 4–16 times larger. For IMSRG, typical jobs now require on the order of 10K–100K core hours, with isolated cases reaching on the order of 1M core hours. Uncertainty quantification will require multiple runs per nucleus, varying basis parameters, truncations, and input interactions. Through 2025, the effort will grow by 3–4 orders of magnitude through the inclusion of the continuum and deformation, as well as the push to larger mass numbers. This takes into account mitigation through more advanced numerical algorithms. GCM

currently uses about 1M hours/year. That will jump by a factor of about 1,000 in 2020–2025 as we include larger single-particle spaces and multi-quasiparticle states.

3.2 Parallelism:

The present codes scale well to the largest machines available today typically using OpenMP and MPI. However, much remains to enable them to efficiently use the dominant new architectures using either many cores per node or GPUs. This effort is essential for the machines coming in the next few years and especially for those appearing as we move to exascale. The QMC codes scale well to from 200K (AFDMC) to 780K cores (GFMC) on MIRA at ANL. They use branching Markov Monte Carlo random walks; GFMC has been particularly tuned to asynchronous dynamic load balancing (ADLB) and distributed memory (DMEM) libraries. The AFDMC case is a bit different because GFMC has much larger memory requirements and OpenMP is more valuable. The GCM codes currently use MPI to distribute calculations of matrix elements of Hamiltonian kernels to many cores. We plan to parallelize the computation of each matrix element via OpenMP. The IMSRG currently employs OpenMP for on-node parallelism and an MPI wrapper for parameter sweeps. For the eigenvalue problems, hybrid OpenMP and MPI parallelism is supported. The use of GPU/many-core architectures to speed up the tensor contractions will be explored.

3.3 Memory:

GFMC presently uses 3 GB of memory per instance of random walk for 12C for energy calculations and 11 GB for neutral current response; approximately 15 GB may be needed for charged current response in $A = 12$ with larger requirements for $A = 14$. The exascale calculations for 16O will need approximately 3 TB for per instance; this will likely require splitting up a configuration across several nodes. We are currently studying the use of high-bandwidth memory on the Intel KNL nodes. We could make good use of “read-only” memory, should that be available on future machines. At present, AFDMC uses much smaller memory per (quasi-independent) random walk, but many more random walks are required. It is also likely that memory per random walk will grow as we move to a hybrid approach that incorporates both AFDMC and GFMC.

IMSRG current evolutions are performed in a shared-memory model due to the coupled nature of the working equations. Currently available systems with 64GB or 128GB are sufficient. Continuum coupling and the explicit treatment of deformation will increase memory requirements in the future by 3 to 4 orders of magnitude, based on the scaling of storage with single-particle basis size. The efficient handling of input $3N$ interactions will become an issue. For current calculations, 20GB–100GB are required to hold the matrix elements in memory, and larger basis sizes and matrices will be necessary to converge observables in ^{136}Xe . The use of memory hierarchies (including high-bandwidth memory) will be required to deal with this issue. For CC, see the “nuclei from first principles” case study. GCM currently uses about 1GB per core to store pieces of two-body density (used as input for IMSRG), and will require 16GB per core for the three-body density.

3.4 Scratch Data and I/O:

In QMC, the random walks are periodically written out for checkpointing and/or for calculating other observables. The data required to restart a calculation of $A = 14$ will be approximately 1 TB. These results could be regenerated at the cost of much larger core-hour requirements. Typical IMSRG calculations require 10 GB scratch space. The checkpoint intervals can be controlled by the user; under current implementation about 5% of the runtime is spent on checkpointing. The checkpoint mechanism itself uses raw binary dumps from memory, which is the most efficient option for the codes at the present scale. More sophisticated parallel I/O will be required in the future. For CC, see the “Nuclei from

First Principles” case study. The GCM code currently writes files on the order of 1 GB for checkpointing occasionally. The size of the files and the frequency of checkpointing will increase dramatically by 2020 as the computations grow in size.

3.5 Long-term and Shared Online Data:

Long-term storage should be available for many observables and for response calculations in QMC. These do not require the full configurations but averages obtained from them for 10–100 configurations apiece. The storage requirements are modest: we expect 20 TB of storage will be sufficient for the next couple of years with requirements eventually moving to on the order of 100 TB. The unitary transformations generated by the IMSRG will be stored to allow the easy evaluation of generic observables. Through 2025, we envision a persistent need for 10–20 TB of active storage for solutions, covering ranges of basis parameters, truncations, and input interactions, as required for proper uncertainty quantification. As input interactions improve, older versions can be retired to archival storage. GCM Hamiltonian and decay-operator kernels are stored, leading to multi-GB files. As we add more HFB states, the active storage will reach 10–100 TB.

3.6 Archival Data Storage:

No specific archival data storage is required other than that under long-term and shared online data. Archival storage of about 20TB should be sufficient through 2025, to ensure reproducibility of intermediate IMSRG results as outdated input interactions are retired.

3.7 Workflows:

QMC workflows are fairly simple, with initial preparation of trial states followed by large-scale calculations checkpointed occasionally. Storage of intermediate configurations is required for calculations of additional observables and history. In IMSRG, separate binaries are chained together by scripts. In the first stage of GCM HFB basis states are generated. These are stored and read in as input for computation of matrix elements.

3.8 Many-Core and/or GPU Readiness:

QMC codes scale very well on MIRA using OpenMP, and we plan to move to theta-like (many-core) machines. The basic structure required is available, but significant work will be required to optimize codes for these new architectures. In IMSRG, the evolution equations and eigenvalue problems are formulated in terms of tensor contractions. In principle, such operations are well suited for implementation on many-core or GPU architectures, but the highly coupled nature of our working equations make data availability (and data locality) as well as communication between host and accelerator highly nontrivial. Dedicated technical support would be extremely helpful for a successful migration. At the lowest level the GCM code repeatedly evaluates matrix Pfaffians. If that procedure could be carried out on GPU,s it would make a great difference. Technical support would be helpful here, too.

3.9 Software Applications, Libraries, and Tools:

The current QMC codes are written in FORTRAN, using FORTRAN 2008 features, and to some degree on BLAS libraries. Also, MPI and OpenMP libraries are used. GFMC uses the ADLB library to scale effectively to the largest machines; AFDMC will move in this direction as well. GFMC also uses the DMEM library for memory management. The current IMSRG code suite uses vendor-optimized linear algebra (BLAS) and MPI libraries, the open-source SUNDIALS suite of ODE solvers, and HDF5 for certain I/O tasks. These will continue to be the cornerstones of the MR-IM-SRG through 2025, and vendor versions are expected to evolve with architectures. Access to vendor-supplied performance analysis tools will be indispensable

for optimizing the codes to specific platforms and architectures. The GCM code uses BLAS and LAPACK. These tools or better versions will continue to be essential.

3.10 HPC Services:

While no definitive needs have been identified at this stage, visualization of the results may become important.

3.11 Additional Needs:

Requirements Summary Worksheet

Codes: All	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	200M	200M	200M
Computational node hours (homogeneous many-core and/or GPU)		400M	2,500M
Memory per node	16 GB	64 GB	64 GB
Aggregate memory	10 TB	100 TB	400 TB
Data read and written per run	100 TB	10X	20X
Maximum I/O bandwidth needed	15 GB/sec	10X	40X
Percentage of runtime for I/O	5%	5%	5%
Scratch file system space needed	30 TB	300 TB	3,000 TB
Permanent online data storage	100 GB	2 PB	5 PB
Archival data storage needed	50 TB	500 TB	2 PB

Quantified Heavy Nuclei

Lead Authors: N. Schunck (Lawrence Livermore National Laboratory) and W. Nazarewicz (Michigan State University)

Contributing Author: S. Wild (Argonne National Laboratory)

1. Description of Research

1.1 Overview and Context:

Our research aims at delivering quantitative predictions for heavy atomic nuclei based on our best knowledge of nuclear forces between nucleons (protons and neutrons) and advanced quantum many-body methods in the framework of nuclear density functional theory (DFT). The Hartree-Fock-Bogoliubov (HFB) equations of nuclear DFT are solved numerically. In the simplest approximation, the single-reference energy density functional (SR-EDF) approach, the DFT equations take the form of nonlinear, dense, medium-size ($\sim 4,000 \times 4,000$) eigenvalue problems. With current energy density functionals, the time to solution is a few hours on a node. To describe complex processes such as fission or $0\nu\beta\beta$, one must compute many (from 1,000+ to 1G+) such SR-EDF configurations, and calculate the quantum correlations between them within the multi-reference energy density functional approach (MR-EDF), which represents a generalized eigenvalue problem.

1.2 Research Objectives for the Next Decade:

The past decade has witnessed the first large-scale SR-EDF applications of DFT, ranging from global surveys of nuclear ground-state properties (masses, limits of stability, etc.) and decay properties (beta decay, fission) to astrophysical applications (simulations of nucleosynthesis, structure of neutron stars). In addition, state-of-the-art statistical methods of uncertainty quantification provided for the first time highly optimized energy functionals and estimates of theoretical uncertainties. As of today, the typical deviations from experiments are on the order of 800 keV for nuclear masses, about 2 orders of magnitude for spontaneous fission half-lives (over a range of more than 35 orders of magnitude), and approximately 30% for fission mass distributions in actinide nuclei. The scientific goal for the next decade is to bring these values down to 400 keV for masses, a factor 2 for fission half-lives, and 10% or less for yield distributions. This will require moving systematically from SR-EDF to MR-EDF calculations together with a 1,000-fold increase in the number of SR-EDF configurations computed (on the order of 1 million to 1 billion for the most complex problems), and the development of more realistic energy functionals, which will include finite-range and/or three-body nuclear forces). All these improvements are expected to increase the cost of computing one configuration by 2 to 3 orders of magnitude.

2. Computational and Data Strategies

2.1 Approach:

We perform large ensemble runs to generate many (thousands now, millions or billions by 2025) SR-EDF configurations characterized by different nuclear shapes, pairing correlations, angular momentum, particle number, etc. Each of these SR-EDF calculations is itself result of a nonlinear eigenvalue problem with a time to solution currently of a few hours per node. Conditions of continuity and information on proximity between neighboring configurations are currently not exploited in our parallel implementations of the ensemble runs. In addition to a static MPI job scheduler, current limitations

include the long time to solution caused by slow and nonoptimized DFT kernels as well as nonscalable, legacy I/O systems.

2.2 Codes and Algorithms:

In the nuclear energy density functional (EDF) approach, the energy density is a functional of the local nucleonic densities, which are obtained by solving the HFB equations. Our DFT solvers **HFBTHO** and **HFODD** use the harmonic oscillator configuration space representation. The expansion coefficients are determined by diagonalization of the HFB matrix, which is defined as the functional derivative of the energy with respect to the densities. Since the HFB matrix depends on the densities, the HFB equations are nonlinear and require an iterative method (self-consistency). Suitable constraints are used to obtain solutions for specific configurations. Large numbers of constraints are handled with a static MPI-based parallel scheduler using both coarse- and fine-grain parallelization. CPU-intensive tasks (construction of densities, calculation of physical observables, and calculation of HFB matrix elements) are accelerated with OpenMP multi-threading. BLAS and LAPACK routines are used for most linear algebra operations.

3. Current and Future HPC Needs

3.1 Computational Hours:

Generating a Cartesian two-dimensional potential energy surface with about 5,000 points currently requires on the order of 1.5M core hours: $5,000 \times (6 \text{ threads}) \times 48\text{h}$. Computational requirements increase when constraints are added. For the most complex problems such as fission, we anticipate that at least five different constraints are needed to reach the scientific goals. Assuming an average of 60–70 points per constraint, this yields about 1G points to be considered. We estimate that a more efficient, dynamical work flow could reduce the number of points effectively computed and recorded to only 1M points (the others being dismissed as nonphysical).

3.2 Parallelism:

Our DFT solver suite uses a hybrid MPI-OpenMP parallel mode with two layers of coarse-grain parallelization. One layer handles the N_x different HFB calculations and a second layer the N_y MPI-ranks needed for each of the HFB calculations. In addition, each MPI-rank is multi-threaded. We intend to optimize on-node execution by implementing accelerators and further spreading the use of OpenMP and implement ADLB scheduler to enable dynamic task allocation at the level of the Masters communicator (this will require an upgrade of the ADLB library).

3.3 Memory:

The minimum shared memory pool is on the order of 16 GB. It would be advantageous to have fast on-chip memory, especially if nodes include large core count (16–128), since one node will be in charge of a single HFB calculation. A long-term goal of 32 GB/node memory, assuming nodes with 128+ cores (possibly connected to a GPU), is desirable. Assuming 16,000 concurrent HFB calculations, this gives a total aggregate memory of about 512 TB.

3.4 Scratch Data and I/O:

We assume 1M converged HFB calculations. Each requires on the order of 200 MB storage (used for checkpointing), which requires 20 TB online scratch storage. Bandwidth per node should be on the order of 200 MB/sec or less. Since not all I/O operations will be concurrent and many HFB calculations will be interrupted and not recorded, it is difficult to aggregate the total bandwidth. I/O should not exceed 1% of the runtime.

3.5 Long-Term and Shared Online Data:

Online storage does not need to be long term (meaning longer than a couple of weeks). Since the DFT community is scattered across the country, it is important to have mechanisms to share the data easily. The goal should be to be able to share 20 TB of data among users of the same machine and among machines for the same user. The transfer of data across machines should be relatively fast.

3.6 Archival Data Storage:

We currently have on the order of 30 TB of data on HPSS tapes. Potential energy surfaces should have a lifetime of about 5 years. Each PES will require on the order of 20 TB storage, and we should anticipate on the order of 5,000 PES (at least one for each nucleus predicted to be stable) on a rolling basis. This means a total of about 100 PB of data at any given time. It will become increasingly important to create a theoretical National Nuclear Data Center (NNDC-t) for the ambitious goals we have for sharing theory simulation results, theoretical interactions, covariances and emulators, wave functions, etc.

3.7 Workflows:

The current work flow is one of the major limiting factors for applications. On facilities such as OLCF's Titan, CPU speed is too slow and runtime is too small. In practice, applications run on smaller systems with faster CPU and longer runtime, such as LLNL's sierra cluster. However, this limits the maximum number of HFB calculations in one run to about 1,024. Because of the lack of dynamic load scheduler, there is no real-time communication between MPI-ranks to optimize the calculation of the PES: HFB calculations that do not converge to the desired solution are not interrupted, resulting in about a 50% loss. By 2020, we should have a dynamic load scheduler implemented in the work flow for the axial DFT solver HFBTHO in order to automatize PES generation and reduce losses. By 2025, the kernels of the symmetry-unrestricted DFT solver should be sufficiently optimized so that the new PES generator can be used efficiently.

3.8 Many-Core and/or GPU Readiness:

We need a dialogue with computer scientists to help us implement accelerators into our codes. To facilitate this task, we will continue to extract critical sections of the code and turn them into small code snippets that can be optimized separately and do not need full knowledge of the code.

3.9 Software Applications, Libraries, and Tools:

We rely extensively on BLAS and LAPACK. Our new work flow will require more scalable I/O software such as ADIOS. Abstractions and tools for performance portability will become increasingly more important as computer architectures evolve. Visualization of the results will become increasingly important.

3.10 HPC Services:

We will rely critically on quality online documentation for compilers (options, optimization), computers (how to submit a job? What software is available? Performance comparisons), accelerators (which ones? How to use them? concrete examples on production machines), libraries. The model for such documentation is NERSC. A dialogue with computer scientists who will work with us to optimize our codes for nonstandard features is indispensable.

3.11 Additional Needs:

We estimate a need for additional 3 domain/computational scientists positions over 5 years in this area. These new positions are critical to enable the effort to scale to the largest scale machines heading toward exascale.

Requirements Summary Worksheet

Code: HFODD	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	10M	1x	1x
Computational node hours (homogeneous many-core or GPU)	-	40M	200M
Memory per node	16 GB	16 GB	32 GB
Aggregate memory	32 TB	64 TB	512 TB
Data read and written per run	1 TB	4 TB	20 TB
Maximum I/O bandwidth needed	0.2 GB/sec	1 GB/sec	4 GB/sec
Percentage of runtime for I/O	1%	1%	1%
Scratch file system space needed	1 TB	10 TB	100 TB
Permanent online data storage	1 TB	10 TB	200 TB
Archival data storage needed	30 TB	100 TB	1,000 TB

Dense Nucleonic Matter

Lead Authors: J. Carlson and S. Gandolfi (Los Alamos National Laboratory)

1. Description of Research

1.1 Overview and Context:

The properties of dense nucleonic matter are critical to the understanding of neutron stars and core-collapse supernovae. To understand these astrophysical objects, one needs to know the equation of state of dense nuclear matter and its response to electroweak probes like neutrinos. Numerical simulations are based on nucleon–nucleon (NN) and three-nucleon (3N) interactions that reproduce the saturation properties of medium and heavy nuclei to enable predictions of nucleonic matter. Electroweak current models are also required to provide a realistic description of the dynamic properties of dense matter, including neutron star cooling and the spectra of different flavors of neutrinos and antineutrinos emitted from the proto-neutron star formed in a core-collapse supernova. These properties are important for efforts in computational nuclear astrophysics, including simulations of neutron stars, their mergers and core-collapse supernovae, and affect supernovae neutrino spectra and the formation of the heavy elements.

1.2 Research Objectives for the Next Decade:

Important new directions in the physics of dense matter are stimulated by recent astrophysical observations. Discoveries of neutron stars as heavy as two solar masses require a stiff equation of state at high densities. Observations of X rays from neutron stars are giving us more information about their radii, and measurements of gravitational waves from merging neutron stars will open new directions in the upcoming years. Anticipated galactic supernovae will tell us more about the interior of core-collapse supernovae and about neutrino physics. Both supernovae and neutron star mergers are considered as possible sites for the astrophysical r-process, which is believed to be responsible for the creation of approximately half of the neutron-rich atomic nuclei heavier than iron. Here, large-scale computations are required to move beyond simple models to accurately predict the structure and dynamics of dense nucleonic matter. Numerous studies have made predictions for the equation of state of neutron matter at zero temperature, particularly below and at nuclear saturation density. However, little is known about high-density behavior. For example, the superfluid gap in low-density neutron matter is now reasonably well understood, but our knowledge of the p-wave neutron and proton superfluidity at higher densities, important for the dynamic properties of neutron stars, is very limited. Studies of dense matter at finite temperature are important for supernovae and neutron star mergers, especially in the context of its response to neutrinos. The initial neutrino spectra at the surface of the proto neutron star play an important role in the explosion mechanism and also in r-process nucleosynthesis. Present-day calculations are based on simplistic assumptions and are not adequate to the comprehensive understanding we require.

2. Computational and Data Strategies

2.1 Approach:

Methods including Auxiliary Field Diffusion Monte Carlo (AFDMC), coupled cluster, and in-medium SRG are required to study dense matter. They have a favorable scaling in particle number, which allows one to study 50–100 nucleons. Other methods including configuration interaction and Green's function Monte Carlo may be able to give valuable tests for small numbers of nucleons, which can be directly compared to lattice QCD calculations. Related methods are being developed to study superfluidity, response, and transport properties.

2.2 Codes and Algorithms:

The methods used to study dense matter are quite similar to those used to study large nuclei and include coupled cluster, in-medium SRG, and AFDMC. Quantum Monte Carlo codes solve the Schrödinger equation by using Markov chain Monte Carlo to sample the path integral used to project out the ground state. These branching random walks employ significant linear algebra at each step of the simulation. These same methods are used to study the structure and response of medium mass nuclei.

3. Current and Future HPC Needs

3.1 Computational Hours:

For a single run, the AFDMC code scales well up to about 130,000 cores on MIRA at Argonne and similarly at NERSC; it yearly uses 5–10 million core hours aggregated between institutional computing provided by Los Alamos and NERSC. We expect our requirement to increase by a factor of 20 by 2020 and by another 1050 by 2025, depending upon available resources. The major requirements driving this increase are the ability to simulate systems with ~10% protons similar to the composition of neutron stars, and the requirements to understand the neutrino response of dense matter as a function of neutron and proton densities and temperature.

3.2 Parallelism:

Currently, AFDMC uses MPI tasks; we have implemented OpenMP as well. The latter is likely to play a more significant role as we move to many-core architectures. It is anticipated that we will change load balancing to ADLB or a similar library that manages quasi-independent tasks.

3.3 Memory:

Memory is currently not a major issue for AFDMC, but the memory needs are likely to grow significantly, as we incorporate more sophisticated algorithms to deal with the fermionic sign problem and the dynamic response. This will require us to explore the use of high-bandwidth memory and burst buffer in new generation machines.

3.4 Scratch Data and I/O:

The AFDMC can easily do checkpoints by writing the configurations at a given imaginary time on the disk. The output for ground-state calculations is usually limited to a few hundred MB for a given ensemble of configurations. At this point the initial and final I/O for a run requires only a few percent of the total run time. This will likely grow significantly with advanced algorithms and response and may require several TB for storage and the use of parallel I/O.

3.5 Long-Term and Shared Online Data:

For large systems, in which the propagation in imaginary time is very time-consuming, we plan to store configurations in such a way that other observables can be post-processed, as commonly done in lattice QCD calculations. Certain data should be stored long term for later analysis and made available online for later analysis.

3.6 Archival Data Storage:

We do not expect to have a particular need for archival storage other than what is described in long-term and shared online data above.

3.7 Work Flows:

An initial data set and list of random walks are prepared initially; the Markov chain Monte Carlo algorithms evolve these systems to larger imaginary time with data analysis running concurrent with the random walks. The final and intermediate states of the simulation are stored; they can be used for analysis of additional observables.

3.8 Many-Core and/or GPU Readiness:

We have implemented MPI with OpenMP for the AFDMC code. At present, the memory and communication requirements are limited enough that we can run efficiently using an MPI process per core. We expect this to change with many-core hardware, to be available shortly. The code is essentially already using a task-based execution model, which together with the modest memory requirements should allow us to run effectively on the many-core architectures. To the extent possible, we will use local tasks to minimize required movement of data.

3.9 Software Applications, Libraries, and Tools:

We rely on highly optimized BLAS and LAPACK libraries and will likely use other libraries for task management. Also, we require compilers that support the current FORTRAN 2003/2008 standards as well as the upcoming FORTRAN 2015 standard. We use and plan to use various standard tools for version control, debugging, profiling, building, regression and unit testing, and bug and feature tracking. We will require tools for optimization on single nodes of the many-core architecture and for debugging and profiling that are efficient and practical to use at exascale.

3.10 HPC Services:

We will have an increasing need for smaller scale systems to perform pre- and postprocessing.

3.11 Additional Needs:

We estimate a need for additional 2 domain/computational scientists positions over 5 years in this area. These new positions are critical to enable the effort to scale to the largest-scale machines heading toward exascale.

Requirements Summary Worksheet

Code: AFDMC	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 2 (As a factor of column 1)
Computational core hours (conventional)	30M	1X	1X
Computational node hours (homogeneous many-core or GPU)	0	200M	2300M
Memory per node	8 GB	16 GB	16 GB
Aggregate memory	10 TB	100 TB	400 TB
Data read and written per run	100 TB	400 TB	1,000 TB
Maximum I/O bandwidth needed	15 GB/sec	100 GB/sec	500 GB/sec
Percentage of runtime for I/O	5	5	5
Scratch file system space needed	20 TB	200 TB	2,000 TB
Permanent online data storage	100 TB	600 TB	2,500 TB
Archival data storage needed	20 TB	400 TB	2,000 TB

D.4 Case Studies Addressing Cold Quantum Chromodynamics

Hadron Structure

Lead Author: K. Orginos (College of William and Mary and Thomas Jefferson National Accelerator Facility)

Contributors: W. Detmold (Massachusetts Institute of Technology), R. Gupta (Los Alamos National Laboratory), B. Joo (Thomas Jefferson National Accelerator Facility), Huey-Wen Lin (Michigan State University), Meifeng Lin (Brookhaven National Laboratory), D. Richards (Thomas Jefferson National Accelerator Facility), S. Syritsyn (Stony Brook University), and F. Winter (Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

The strong force is one of the fundamental forces of nature giving rise to a rich phenomenology that has been explored for decades through intense experimental and theoretical efforts. However, it is only now that direct connections between the properties of fundamental particles such as protons and neutrons (nucleons) can be computed from the underlying theory of strong interactions: quantum chromodynamics (QCD). This connection has become solid through numerical computations using the lattice formulation of QCD. This approach, which has been available since the discovery of QCD, relies on state-of-the-art supercomputers and only recently have these supercomputers become powerful enough to perform calculations of the simplest observables, such as the mass and the distribution of charge and magnetism in the nucleon. Furthermore, calculations not possible before are becoming accessible now. Computations of generalized parton distribution functions (GPDs) and transverse-momentum dependent distribution functions (TMDs) allow us to form a 3D picture of the nucleon. Furthermore, precise computations of couplings of the nucleon to the weak force (weak matrix elements) will aid experimental searches for physics beyond the standard model of particle physics (BSM physics). With high-performance computing (HPC) entering the exascale era over the next decade, these calculations will reach unprecedented precision, enabling scientific discovery through comparison with experiments and dramatically improving our understanding of the subatomic world.

1.2 Research Objectives for the Next Decade:

Hadron structure calculations aim to precisely determine fundamental quantities characterizing the nucleon, including form factors, moments of parton density, helicity, and transversity distributions, moments of generalized parton distributions (GPDs) and transverse-momentum dependent distribution functions (TMDs) using the fundamental theory of QCD. These observables are directly relevant to the experimental programs at Jefferson Lab 12 GeV, RHIC-spin, and Fermilab and at the future electron-ion collider (EIC). Physical observables are computed through correlation functions that are evaluated in a Euclidean field theory using Monte Carlo. The computational challenge at the exascale is to perform computations at a very fine discretization scale (lattice spacing), so that reliable extrapolations to the continuum limit can be performed. In addition, given that the numerical computation is performed in a

finite volume, sufficiently large volumes have to be used in order to control the infinite volume extrapolation. Finally, in order to achieve our precision goals, calculations with physical parameters have to be performed. At the energy regime relevant to hadron structure, there are four input parameters in QCD that are set by matching onto observed particle masses: the three light-quark masses (up, down, and strange) and the dynamically generated scale of strong interactions, Λ_{qcd} , which determines the energy dependence of the interaction strength. Once this is done and after continuum and infinite volume extrapolations are performed, true *ab initio* predictions of QCD are achieved with quantifiable and controlled uncertainties without any model dependence.

2. Computational and Data Strategies

2.1 Approach:

The central computational task of lattice QCD calculations is a Monte-Carlo evaluation of the multi-dimensional integrals that define the correlation functions from which physical observables are extracted. As such, they proceed in two phases: generation of ensembles of gauge fields and the calculation of correlation functions on each gauge configuration. Gauge configuration ensembles are generated via a Markov chain process: the Hybrid Molecular Dynamics Monte Carlo method (HMC). On each gauge configuration, correlation functions are constructed from quark propagators that are computed by solving a sparse system of linear equations (the Dirac equation). Modern methods of constructing correlation functions require approximately 10,000 linear system solves per configuration. However, correlation functions on each configuration can be evaluated independently, resulting in extra parallelism. The success of these methods has been enabled by a combination of hardware and algorithmic advances (adaptive multi grid). This step of the calculation requires a large scratch space to save data, that is, short-term storage. Future calculations can make use of burst buffers and data partitions to offload I/O from the main compute processes. The final stage of constructing correlation functions involves tensor contractions. Depending on the physics goals and technique used, these can have a very high cost in FLOPS and also be bottlenecked by I/O to some extent. This stage of the calculation can generate thousands of correlation functions per configuration. Currently, these are stored in file-based databases. However, future calculations may make use of burst buffers and database technologies for the most efficient handling of these data.

2.2 Codes and Algorithms:

There are several openly available LQCD code bases in the United States, primarily developed and maintained by the U.S. Lattice Quantum Chromodynamics (USQCD) Collaboration through SciDAC funding. The main workhorse code suite for our project is the Chroma software suite. In addition, correlation function construction utilizes RedStar. Ensemble generation is done using Hybrid Monte Carlo (HMC), in which a major computational task is the solution of linear systems of equations. The quark propagators that are used to calculate correlation functions are also obtained through the solution of linear systems. Our codes contain optimized linear system solvers that include Krylov and AMG solvers. Each correlation function computation is a task that involves tensor contractions. For hadron structure calculations, the tensors involved are rank 3 and rank 2. Current contraction algorithms are optimized to run on various platforms including GPUs and Xeon Phi.

3. Current and Future HPC Needs

3.1 Computational Hours:

The hadron structure project currently is consuming approximately 3M GPU hours a year in gauge generation on various platforms. For the computation of matrix elements, last year's allocations were about 10M GPU hours on both leadership-class facilities and USQCD resources. From now until 2020, the project requires gauge ensembles with a range of masses close to the physical point, a range of space-time volumes, and at least three lattice spacings in order to take the continuum limit. The total computational resources required for these calculations are estimated to be 10B GPU hours for ensembles with 1,000 configurations each. This includes three volumes (5, 6, and 8 fm), three pion masses (140, 180, and 200 MeV) and three lattice spacings (0.045, 0.06, and 0.077 fm). Between 2020 and 2025, electromagnetic effects will be included in the gauge ensembles. A comprehensive study of these effects is estimated to require 100B GPU hours for similar ranges of parameters as for the ensembles without electromagnetism. The matrix element calculations will be performed with a new methodology that has been shown to be successful for spectroscopic calculations. This approach allows for amortization of the computational cost because intermediate results can be reused for other projects as well as novel calculations that are not identified. For this reason it is valuable to store these intermediate results. Currently, 16B GPU hours are required to process the configurations generated through 2020. To include electromagnetic effects, approximately 64B GPU hours are required. Both components of the project require a large amount of computation that can be accomplished only in a multi-year effort. In the table, we present the computational requests for years 2020 and 2025. In 2020 we envision generating configurations with lattice spacing 0.06 fm at the physical pion mass point in an 8-fm spatial volume and performing calculations of correlation functions on configurations with a spatial volume of 6 fm. The same calculations including electromagnetic effects will be performed during 2025.

3.2 Parallelism:

Ensemble generation scales well up to a large fraction of today's leadership-class supercomputers. In the future, similarly large fractions of the leadership machines will be required in order to generate the ensembles described above. The analysis phase of the project uses throughput computation of the Dirac-equation solution vectors for many right-hand sides on numerous, but small, partitions of the system. It can use either accelerators or many-core systems. The large number of solution vectors is partially contracted to form tensors for each time slice of the lattice and must either be paged to disk or held in memory. The contraction that consumes these tensors typically uses single many-core or GPU nodes for each time slice. The Hadron node code relies on batching subtask to optimize the tensor contractions. Finally, each configuration worth of computations can be executed in parallel in such a throughput mode.

3.3 Memory:

Memory is more of an issue for the analysis phase than the gauge generation phase. The Dirac solvers are intentionally run in small partitions to increase their efficiency. The limit of this scaling is the memory size per node. Today, GPU and KNL nodes are currently used with up to 128 GB of main memory. The solver codes use specialized caching routines to page data in/out of fast-memory to slower main memory, and we expect this technique will remain essential in the future.

The contraction step is also typically memory-limited and requires the most care in the work flow. The tensors, when in a dense representation, are expected to occupy approximately 11 GB on a single computational node for the anticipated 6-fm lattice sizes. Future developments will switch to a sparse (banded) structure for which approximately 1 GB is required. There are numerous such structures, not

all of which are expected to fit in memory. The code will have to page these out to scratch space in order to fit into memory.

3.4 Scratch Data and I/O:

The main challenge for the analysis phase is the large I/O demands. Such computations are carried out independently for each of the approximately 1,000 configurations. A large number of Dirac equation solution vectors must be processed. In the second phase of the work flow, the solution vectors are partially contracted and saved on disk in a form that allows them to be reused in subsequent contractions. These objects, called “perambulators,” are rank 2 tensors and will be stored in long-term storage in databases that facilitate random access and retrieval. For the 6-fm lattices, the perambulators are 34 TB per configuration. In addition, higher order tensors are needed that can be cached and re-used in a similar fashion. While the tensor contractions proceed, there will be intermediate temporaries that will be created and reused, and this will reduce the overall computational cost. Typically, approximately 100 GB in temporaries is anticipated. The I/O bandwidth required to write and read these objects is expected to be large enough that rapid paging to burst-buffers will be required. Our present work flow optimizes wall-clock time by balancing I/O versus computation.

3.5 Long-Term and Shared Online Data:

Long-term storage for the propagators and the generalized propagators (that include the operators needed for the matrix elements) is required. For 1,000 configurations and all the ensembles, this amounts to about 250 PB of long-term storage. If all the solution vectors for just one source time-slice were stored, the maximum (i.e., largest box and smallest lattice spacing) I/O demands would be 10 TB per configuration. However, given the amount of data, computations would be coordinated to use the solution vectors as they are constructed and discard them once complete. Going beyond 2020 and including E&M effects, the long-term storage requirements will grow by an approximate factor of 4, while the short-term storage needs will remain approximately the same.

3.6 Archival Data Storage:

A total of 250 PB of archival storage will be required for long-term re-use of data. Calculations including E&M effects will require approximately four times that amount.

3.7 Work Flows:

Currently the coordination for the analysis phase involves mostly scripts and C++ codes for the computationally intense parts. We envision developing a more comprehensive work flow code suite that accommodates a greater portion of our community from HEP and Nuclear Physics.

3.8 Many-Core and/or GPU Readiness:

Our software supports many-core and GPU systems. The Dirac solvers have been developed under SciDAC and in collaboration with Nvidia and Intel. We are actively extending our basic data-parallel software framework used by Chroma and QDP++ as a part of SciDAC. The tensor contractions performed in our Hadron package also support many-core and GPU systems.

3.9 Software Applications, Libraries, and Tools:

Our codes rely on the USQCD-developed software for lattice QCD calculations. It is based on C/ C++ codes that rely on MPI for communications.

3.10 HPC Services:

A database system for archiving the most important data of our calculations will be required. This will ensure easy access to the data we generate for the members of our large and diverse collaboration.

3.11 Additional Needs:

In order to take full advantage of the future high-performance supercomputers, it is essential that smaller scale systems be deployed within universities where the training of graduate students takes place. With such resources, new ideas can be explored and the next generation of computational scientists can be trained to effectively use the leadership-class machines that will be available to our projects in the future.

Requirements Summary Worksheet

Code: Chroma and RedStar	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)*	100M		
Computational node hours (homogeneous many-core)**	2M	Can use either CPU or GPU	Can use either CPU or GPU
Computational node hours (with GPU or accelerator)***	2M	350M	3.3B
Memory per node	128 GB	1x	1x
Aggregate memory	8 TB	1x	1x
Data read and written per run	1 TB	3 TB	10 TB
Maximum I/O bandwidth needed	1 GB/sec	30 GB/sec	30 GB/sec
Percent of runtime for I/O	10	20	20
Scratch file system space needed	10 TB	470 TB	800 TB
Permanent online data storage	25 TB	100 TB	1 PB
Archival data storage needed	100 TB	250 PB	1,000 PB

* Core hours are used for conventional" processors (i.e., node hours × cores per node). Intel Ivy Bridge is an example.

** Node hours are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi Knights Landing is an example.

*** Node hours are used for GPU or accelerator usage.

**** For example, 32 times column 1.

Hadron Spectroscopy

Lead Authors: Robert Edwards, Balint Joo, and David Richards (Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

Spectroscopy has played, and continues to play, a pivotal role in the development of our understanding of the strong interaction. Recent experiments worldwide have upended the conventional picture that hadrons, the basic building blocks of matter, have a simple classification scheme, and the observation of these intrinsically new hadrons has ignited a firestorm of interest to reclassify the spectrum. Quantum chromodynamics (QCD) is the theory that should describe this spectrum, but it also suggests new forms of matter that have not been clearly observed. Lattice QCD is the only known systematic way to calculate the spectrum in terms of quark and gluon degrees of freedom. This project addresses a unique challenge for our understanding of QCD: How does the complexity of the hadron spectrum emerge from these degrees of freedom? In particular, a key goal is to identify the presence of an exotic state of matter called “hybrids”—states admitted by the QCD in which the gluonic degrees of freedom are manifest and will be searched for in the new GlueX experiment at Jefferson Lab and other facilities worldwide.

1.2 Research Objectives for the Next Decade:

A key goal of this project is the determination of resonance properties of the hadron spectrum. The restriction of the calculations to a finite spatial volume is a tool that permits indirect access to map the energy dependence of scattering amplitudes from which resonance properties can be determined. The predictions from these calculations will guide experimental searches at GlueX and confront future predictions in other sectors of QCD. Notably, in the charmonium sector of QCD, future calculations will address the origin of new experimentally observed, and quite unexpected, resonant structures that require a new understanding of the dynamical degrees of freedom of quarks and gluons. Future calculations will also address precision determinations in these systems in which electromagnetism is important. The challenge is the computation of the highly excited, finite-volume, energy spectrum directly from QCD, on many lattice sizes and many moving (spatial-momentum) frames, to allow for a precise mapping of the energy dependence of these scattering amplitudes. For this, highly scalable source construction techniques are needed that allow for the calculations to be carried out on very large lattices, and a new software and workflow infrastructure that can support the intense I/O demands placed on the computing systems.

2. Computational and Data Strategies

2.1 Approach:

The three phases of the work flow involve generating lattice-gauge configurations, the solution of the lattice Dirac equation on each configuration for many different source constructions, or right-hand sides, and the contraction of these solution vectors to form Euclidean correlation functions. From these functions the excited energies, the main outcome of our calculations can be determined. Key to the method is the construction of matrices of these functions featuring a large basis of many-body operators that resemble the scattering states expected in the spectrum. The contraction of the quark fields

requires solution vectors on every time slice of the lattice and for many source constructions on each time slice. This is a challenging computational task that requires the development of architecturally optimized Dirac solvers supporting nearly 1M right-hand sides. The contraction of the solution vectors requires a carefully coordinated work flow for the creation of tensors representing the operators and the evaluation of the tensor contractions. Intermediate paging of these tensors to disk requires investigations into new database and I/O technologies. Future algorithm development will focus on different evaluation strategies that are more optimal for many-body baryon operator constructions.

2.2 Codes and Algorithms:

The Chroma code is used for the first two of three phases of the work flow: gauge generation and Dirac equation solvers for GPUs and many-core CPUs. The Chroma code also supports the creation of intermediate data that are used in the third stage of the work flow, namely, the contractions. The construction of the correlation functions is implemented using the “distillation” method, and the Redstar code was developed to support many-body operator constructions and automate the optimal edge-reduction of graphs for the tensor contractions. The Redstar code is built over a “Hadron” API that supports tensor contractions with implementations that utilize threaded, but single-node, ZGEMM routines or hand-coded routines for performance on various architectures, such as CPUs, GPUs, and KNLs. The Hadron code relies heavily on batching for performance.

3. Current and Future HPC Needs

3.1 Computational Hours:

The first phase of the spectroscopy project in the 2020 time frame will focus on the determination of the exotic meson spectrum in the physical limit of QCD and the charmonium spectrum, where E&M is not necessary. The computations envisioned will use first the dynamical Clover gauge fields with 6-fm spatial extent and lattice spacing 0.06 fm. Resolving high-momentum states is crucial in the spectroscopy campaign, and this will be achieved by using a large number of source constructions for the Dirac solver. While the contractions are smaller in computational cost, the main challenge is satisfying the strenuous I/O demands. The second phase of the project for the 2025 time frame will focus on extending these calculations to include E&M, which is needed in order to investigate the low-energy sector of QCD involving decays of mesons, where the effects of isospin breaking are crucial, and near threshold states in charmonium.

3.2 Parallelism:

The analysis phase of the project, described here, uses throughput computation of the Dirac-equation solution vectors for many right-hand sides on numerous, but small, partitions of the system. It can use either accelerators or many-core systems. The large number of solution vectors is partially contracted to form tensors for each time slice of the lattice and must either be paged to disk or held in memory. The contraction step, which consumes the tensors in memory or on disk, typically uses single many-core or GPU nodes for each time slice of the contractions. The Hadron node code relies on batching to optimize the tensor contractions. Here, each spin component of the tensor contraction is passed to a smaller part of the node to increase memory and performance within that smaller partition.

3.3 Memory:

Memory is an issue for the analysis phase where the Dirac solvers are intentionally run in small partitions to increase their efficiency. The limit of this scaling is the memory size per node. GPU and KNL nodes with up to 128 GB of main memory are currently used. The solver codes today use specialized

caching routines to page data in/out of fast memory to slower main memory, and this technique is expected to be essential in the future. The contraction step is also typically memory limited. Future developments will switch to a sparse (banded) structure to save memory and contraction costs.

3.4 Scratch Data and I/O:

In the second phase of the work flow, large numbers of Dirac solution vectors are partially contracted into objects called “perambulators” and stored in databases as well as tensor objects representing the operators. During computations, there will be about a few hundred intermediate temporaries, about 1 GB each, to be created per time slice (computational node). The I/O bandwidth required to write and read these objects is expected to be large enough that some kind of rapid paging to burst buffers or network database will be needed. The trade-off for the computations is the paging of these temporaries versus recomputing them to minimize the overall wall-clock time.

3.5 Long-Term and Shared Online Data:

The long-term storage for the perambulators is required, and for 1,000 configurations, this amounts to about 30 PB each for each of the two lattice sets. If all the solution vectors for one source time slice were stored, the I/O demands would be 250 PB and 800 PB over the period. However, given the amount of data, computations would be coordinated to use the solution vectors as they are constructed and discard them once complete.

3.6 Archival Data Storage:

The total storage used by the spectroscopy project, to date, is about 1.5 PB. We anticipate requiring approximately 100 PB of archival storage over the course of the project.

3.7 Work Flows:

Currently, the coordination for the analysis phase involves mostly scripts and C++ codes for the computationally intense parts. We envision developing a more comprehensive work-flow code suite that accommodates a greater portion of the HEP and Nuclear Physics community.

3.8 Many-Core and/or GPU Readiness:

Extensive software development has been under way to support many-core systems and GPU systems. The Dirac solvers is one area of active development under SciDAC and in collaboration with Nvidia and Intel. We are actively extending our basic data-parallel software framework used by Chroma and QDP++ as a part of SciDAC. The other area is support for tensor contractions within our Hadron package. We expect to take advantage of these systems in the future.

3.9 Software Applications, Libraries, and Tools:

We rely heavily on the software stack developed under SciDAC, and seek continued collaborations with industry and also lab partners. There is a continued need for BLAS routines, such as ZGEMM, optimized for future architectures. C++ support is required.

3.10 HPC Services:

We currently use database software derived from Berkeley Sleepy Cat for the storage of intermediate data. We are interested in investigating new technologies, such as network-based storage solutions, which may have better accessibility on compute nodes.

3.11 Additional Needs:

Our projects rely heavily on training provided to graduate students and postdoctoral fellows at our host institutions, for example, as a part of SciDAC. These junior scientists are trained to develop software and to work and use HPC facilities. We also are crucially dependent on local compute resources. These resources have been a critical part of the success of USQCD. We have relied heavily on the leadership resources for our most demanding computations and the local compute facilities for the continued analysis of these data. We rely on this mutually beneficial coordination to continue into the future.

Requirements Summary Worksheet

Code: Chroma and Redstar	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)*	100M		
Computational node hours (homogeneous many-core)**	2M	Can use either CPU or GPU	Can use either CPU or GPU
Computational node hours (with GPU or accelerator)***	10M	Gauge: 101M Props: 313M Contract: 30M	Gauge: 1015M Props: 616M Contract: 40M
Memory per node	128 GB	1x	1x
Aggregate memory	8 TB	100 TB	400 TB
Data read and written per run	3 TB	250 TB	800 TB
Maximum I/O bandwidth needed	1 GB/sec	100 GB/sec	100 GB/sec
Percentage of runtime for I/O	10	20	20
Scratch file system space needed	10 TB	500 TB	1,000 TB
Permanent online data storage	200 TB	500 TB	1,000 TB
Archival data storage needed	1,000 TB	5,000 TB	5,000 TB

* Core hours are used for conventional processors. (i.e., node hours × cores per node). Intel Ivy Bridge is an example.

** Node hours are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi Knights Landing is an example.

*** Node hours are used for GPU or accelerator usage.

**** For example, 32 times column 1.

Nuclear Physics from Lattice QCD

Lead Authors: William Detmold (Massachusetts Institute of Technology) and Martin Savage (Institute for Nuclear Theory)

Contributors: Balint Joo (Thomas Jefferson National Accelerator Facility), Kostas Orginos (College of William and Mary and Thomas Jefferson National Accelerator Facility) and Frank Winter (Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

Lattice QCD calculations play a critical role in refining the nuclear forces and interactions that determine the nature of (the visible) matter. Providing unique input into nuclear many-body calculations, these forces are essential in predicting the structure and decays of nuclei, nuclear reaction rates, the response of nuclei in electroweak processes, and the behavior of matter in extreme astrophysical environments. Many quantities can be determined that are difficult (or impossible) to access experimentally, thereby providing theoretical support that is critical in capitalizing on the investments made in the U.S. experimental program. The forces that describe the stable and long-lived nuclei have been tightly constrained through decades of experiments. However, the uncertainties in the nuclear forces that determine the structure and behavior of neutron-rich nuclei and of the matter that forms during the core-collapse of supernovae are limiting the precision of predictions for the structure and dynamics of these systems. The Facility for Rare Isotope Beams (FRIB) experimental program will reduce some of these uncertainties through direct measurements of processes involving short-lived, neutron-rich nuclei. To complement this effort, lattice QCD calculations aim to refine other components of the forces through studies of neutron-rich systems and hypernuclei that are not accessible in experiment.




Figure 1. The magnetic moments and polarizabilities of light nuclei are now being calculated with lattice QCD. (Image reproduced with permission from William Detmold.)

1.2 Research Objectives for the Next Decade:

The spectrum, properties, and interactions of light nuclei, multi-nucleon and neutron-rich systems, as well as hypernuclei will be calculated at and near the physical quark masses, including electromagnetism, leading to a significant refinement of the nuclear forces. The results of these studies, through collaboration with nuclear many-body theorists, will be incorporated into nuclear many-body calculations of quantities of importance in nuclear physics to support and complement the experimental nuclear and high-energy physics programs planned for JLab, FRIB, the EIC, LBNF/DUNE and in searches for double-beta ($0\nu\beta\beta$) decay.

These calculations will permit dissections of the nuclear forces at the levels required for precision calculations in nuclei and dense matter. They will be performed over a range of lattice spacings and in multiple space-time volumes to enable extrapolations to infinite volume and the continuum. In addition to refining the multi-body interactions between protons and neutrons, this extensive series of calculations will tightly constrain the three- and four-neutron interactions and provide precision hyperon-nucleon forces, quantities that are vital in predicting the behavior of dense matter, but have so far proven elusive to laboratory experiments. The responses of nuclei to electroweak probes and candidate dark matter particles are key quantities in the analysis of neutrino and dark matter direct detection experiments; the multi-body aspects of these interactions will be precisely determined. Achieving these goals requires continuous advances in all levels of HPC, storage, networking, algorithms, and work flows.

2. Computational and Data Strategies

2.1 Approach:

Lattice QCD is a numerical technique in which space-time is discretized, and the path-integral dictating the nonlinear quantum dynamics of the quark and gluon fields is evaluated by Monte Carlo techniques. Ensembles of field configurations with a range of lattice spacings and volumes are required at each selected set of light-quark masses in order to extrapolate to the continuum and infinite volume. Resources requirements increase with decreasing lattice spacing and increasing volume. Each configuration is saved to disk and archived for use in subsequent calculations. Large-scale configuration generation occurs over a period of months to years and requires *capability* HPC. A large number of light-quark propagators and low-lying eigenvectors are determined on each configuration through solution of large sparse linear systems. Algebraic multi-grid and deflation of low-mode eigenvectors are used for this task. On a single-field configuration, the intermediate storage can reach 20 to 35 TB, and performant I/O becomes crucial on the large problem sizes. Propagators typically require HPC *capacity* resources; however, propagators in large volumes and small lattice spacing can require *capability* resources. Correlation functions, produced from propagators and hadronic blocks (intermediate stages of calculations) are written to disk, archived, and moved off-site for subsequent analysis that generally can be performed on modest-capacity resources. Their generation requires HPC *capacity* facilities but can be performed equally well on *capability* resources.

2.2 Codes and Algorithms:

The Chroma lattice field theory code suite was designed as part of the DOE SciDAC effort within the USQCD initiative and has evolved to meet our physics goals and to accommodate the multi-pronged evolution in computer architectures. Built over a C++ data parallel API and DSL implementation called QDP++, it provides an architecture independent programming API along with I/O support and is designed following modern software engineering programming practices. It can be compiled in several architectural modes supporting a hybrid communications and threads model. The parallel version is built over a USQCD communications package called QMP, which wraps MPI as well as other architecturally specific hardware communication variants. Further, QDP++ has been implemented using a just-in-time compilation framework that generates and executes code on GPUs. Using a hybrid MPI/OpenMP model, linked against the QUDA GPU solver library, this code is in production utilizing more than 4,000 GPU nodes.

The **NPLQCD** code provides a rapid evaluation of the quark contractions required to generate nuclear correlation functions by automated code generation from recursion techniques on tensors. An optimized basis of nuclear sources and sinks is used to generate correlation functions to be subsequently processed using advanced statistical analysis techniques on local capacity resources.

3. Current and Future HPC Needs

3.1 Computational Hours:

In 2015–2016, nuclear spectroscopy and structure and interactions calculations have received the equivalent of 200M Titan core hours (ALCC 57 M, NERSC 47 M, USQCD 105 M) using Titan, Edison, Cori, and USQCD-capacity GPU and CPU clusters. Significant growth in the amount of time required for future calculations is anticipated. In recent years, the viability of a QCD approach to nuclei has been demonstrated, and investigations of the basic features of nuclei and the associated scalings of a number of novel algorithms developed for these calculations have been performed. There are a number of extensions that will be pursued in the coming decade that expand the cost of the required calculations. To estimate the required computational resources, various projections of future scenarios have been investigated. The results of these projections are shown in the table.

3.2 Parallelism:

Configuration generation is accomplished with capability resources (typically in the thousands of GPUs or tens of thousands of CPU cores), followed by a very large number of propagator and measurement tasks, each of which can be performed on smaller partitions (although some techniques require memory footprints available only in capability systems). The integrated wall time spent on measurements now exceeds that needed for configuration generation. Heterogeneous machines such as Titan, as well as more conventional machines such as Edison and Cori phase I and USQCD clusters, are currently utilized. On Titan, automatic code generators use the QDP-JIT software stack to run efficiently at scales of up to 5,000 GPUs. CPU and GPU resources are used concurrently. On Edison, typical measurement jobs are in the 4,000–8,000 core range. Without new algorithms, this style of production is expected to continue in the exascale era. As multi-scale methods mature, the configuration generation phase is expected to become amenable to task-parallelism, exploiting several architectural features concurrently. Estimates of the parallelism exhibited by calculations in 2020 and 2025 result from naively scaling the current production requirements, assuming the space-time volume held locally per processing element is fixed at its current value. Both the generation of gluon configurations, propagators, and contractions are expected to run on 10,000–100,000 future accelerators.

3.3 Memory:

Currently many GPUs are used in parallel in a single run with the optimum footprint that uses most of the memory on each GPU. This is anticipated to remain the case going through the exascale era as much of the calculation can be massively parallelized through working on many configurations simultaneously.

3.4 Scratch Data and I/O:

Assuming future machines maintain mean-time-between-failure rates that allow runs extending over multiple hours, it is natural to checkpoint at configurations where measurements will be performed. If not, more resilient software, which is aware of node failures and can readjust to them, will be necessary. Given that I/O is not expected to increase as fast as computational capacity, other scenarios do not appear viable.

3.5 Long-Term and Shared Online Data:

Storing gluon configurations is essential as they form a community-wide resource, which are expected to require approximately 24 PB (4 PB by 2020 and another 20 PB by 2025). It is impractical to keep all of the generated propagators, although subsets of them will be used in downstream analyses.

3.6 Archival Data Storage:

A large amount of data is archived at Jefferson Lab, NERSC, and Oak Ridge: in total approximately 0.4 PB of current and past calculation results. Configurations and a small subset of other data structures are expected to be stored in the exascale era.

3.7 Work Flows:

Current work flows consist of three phases: configuration generation, propagator calculations, and contractions that yield the nuclear correlation functions, with the first two phases common to a range of calculations in nuclear and high-energy physics. The physics of nuclei and nuclear systems represents one of the central drivers for the isotropic clover gauge generation program that is currently being undertaken, both in terms of choice of simulation parameters and in terms of workforce effort. This is expected to continue as a community project, with other scientists in the community joining the effort. Given ensembles of field configurations, the extraction of physics for multi-nucleon systems proceeds through the calculation of quark propagators that are subsequently combined. The study of matrix elements in nuclei requires additional solutions of the Dirac equation. For many nucleon systems, a number of sophisticated contraction techniques have been developed. These make use of various BLAS and Lapack routines.

3.8 Many-Core and/or GPU Readiness:

The computationally intensive aspects of configuration generation and propagator calculation are performed using libraries that are highly optimized for the currently available hardware. These libraries have been developed in close coordination with engineers at Nvidia and Intel, and we expect this collaboration to continue and expand in the exascale era. The technology for performing the quark contractions used to generate nuclear correlation functions is currently transitioning. The present method has been implemented efficiently on GPUs using the QDP-JIT framework and will be optimized for MIC within the year. On the 2025 time scale, we expect new approaches to these many-nucleon problems will be developed and the features of new machine architectures will strongly influence the software design.

3.9 Software Applications, Libraries, and Tools:

Studies of nuclear interactions and structure make use of codes that extend the Chroma field theory library and the USQCD SciDAC software stack (qmp/qio/qdp/qdp++). Propagator calculations make use of the QUDA inverter library on GPUs and the QPhiX library on MIC. Software for the computation of many-body contractions uses BLAS and Lapack routines, benefitting from optimized vendor libraries.

3.10 HPC Services:

While no definitive needs have been identified at this stage, visualization of the results of lattice QCD calculations may become important. Field configurations will be centrally archived for the community to use, subject to sharing agreements; it may be beneficial for such repositories to become more standard with automatic and secure provenance management.

3.11 Additional Needs:

Lattice QCD calculations are workforce-intensive, requiring HPC-trained teams of scientists with diverse skill sets. A number of postdoctoral fellows, laboratory scientists, and research faculty at universities dedicated to the execution of lattice QCD calculations are currently supported by the USQCD SciDAC project, including those contributing to developing the nuclear forces component. Support of this nature is essential through the exascale era. Local computational resources at PI institutions and those operated by USQCD at Jefferson Lab, Brookhaven, and Fermilab, have proven invaluable for the rapid turnaround and testing of new ideas and algorithms, circumventing the need for the preparation and approval of computing resource proposals (whose turnaround time is typically months). Such machines continue to play a crucial role in training junior scientists and attracting them to the field.

Requirements Summary Worksheet

Code: Chroma/nplqcd	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)* (nodes composed of conventional CPUs only)	80M (NERSC, Jefferson Lab)	Can use many- core, CPU, or GPU	Can use many- core, CPU, or GPU
Computational node hours (homogeneous many-core)**	0	Can use many- core, CPU, or GPU	Can use many- core, CPU, or GPU
Computational node hours (with GPU or accelerator)*** (single K20x node hours, e.g., Titan node hours) (resources for a given year, <i>not</i> integrated)	5M	x 24 (120M)	x 192 (950M)
Memory per node (per GPU)	8 GB	32 GB	64 GB
Aggregate memory	5 TB	50 TB	200 TB
Data read and written per run	7 TB	100 TB	400 TB
Maximum I/O bandwidth needed	1 GB/sec	100 GB/sec	100 GB/sec
Percentage of runtime for I/O	~10%	~10%	~10%
Scratch file system space needed	100 TB	1,000 TB	5,000 TB
Permanent online data storage	40 TB	500 TB	2,000 TB
Archival data storage needed	500 TB	4,000 TB	24,000 TB

* Core hours are used for conventional processors. (i.e., node hours × cores per node). Intel Ivy Bridge is an example.

** Node hours are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi Knights Landing is an example.

*** Node hours are used for GPU or accelerator usage.

**** For example, 32 times column 1

Fundamental Symmetries

Lead Authors: R. Gupta (Los Alamos National Laboratory), T. Kurth (Lawrence Berkeley National Laboratory and National Energy Research Scientific Computing Center) and S. Syritsyn (Stony Brook University)

Contributors: Balint Joo (Thomas Jefferson National Accelerator Facility), Kostas Orginos (College of William and Mary and Thomas Jefferson National Accelerator Facility) and Martin Savage (Institute for Nuclear Theory)

1. Description of Research

1.1 Overview and Context:

The accepted Standard Model of fundamental particles and their interactions preserves the total lepton-number, L , and baryon-number, B , at the classical level. However, for the observed excess of baryonic matter over antimatter to be generated in the early Universe, charge-parity (CP) and baryon number have to be violated with couplings much stronger than in the Standard Model. For this reason, a number of current and planned experiments explore CP, B , and L -violating processes. In particular, the new US-based DUNE experiment expected to start operation in 2024 will look for B -violating proton decay; novel neutron-antineutron oscillation experiments are also envisaged. An experiment to measure the neutron Electric Dipole Moment (nEDM), one of the most sensitive probes of time-reversal violation, is under development at the Spallation Neutron Source (SNS) at Oak Ridge. Finally, the new flagship initiative in the NSAC Long Range Plan for Nuclear Science is the development and deployment of a ton-scale experiment to search for double-beta ($0\nu\beta\beta$) events that violate lepton number.

These experiments will look for effects of symmetry-violating interactions of quarks and gluons on the properties of nucleons and nuclei. In order to constrain models of new physics, QCD has to predict, with control over all systematics, the magnitude of the nEDM and the rates of proton and $0\nu\beta\beta$ decays induced by CP, B , and L violations in these models. Strong interactions of quarks and gluons within protons, neutrons, and nuclei make analytic calculations intractable. The only available first-principles framework is lattice QCD, whose calculations are extremely computationally intensive. However, the progress in methodology, algorithms, and raw computer power shows that over the next decade lattice QCD can achieve the precision required to fully realize the power of the experiments mentioned above to constrain new physics.

1.2 Research Objectives for the Next Decade:

- **nEDM and CP violation.** Generic new sources of CP violation at the TeV scale manifest themselves at the hadronic scale through a small number of local CP-odd operators, weighted by effective couplings that encode information of the new physics. The CPT invariance of Lorentz-invariant local quantum field theories ensures that CP-violating operators produce T -odd observables. The goal is to compute the matrix elements of the leading (dimension-5 and 6) CP-odd operators in the nucleon. These operators include the light-quark EDMs, chromo EDMs, the Weinberg three-gluon operator, and four-quark operators. Phenomenological studies of models containing new physics have shown that to maximize the constraining power of neutron EDMs, the minimal precision goals are as follows: 50% on the strange quark EDM matrix elements, 25% on the quark chromo EDMs, and 50% on the Weinberg and four-quark operators.

- **Proton decay, neutron-antineutron oscillations, and B violation.** The goal is to eliminate any remaining uncertainties in the matrix elements for these processes. In particular, the new DUNE experiment will be sensitive to the proton decay $p \geq \nu K$. Resources within next decade will be sufficient for precision greater than 10%.
- **Neutrinoless $0\nu\beta\beta$ decay.** The scientific goal is to determine the few-nucleon matrix elements needed to understand the dependence of $0\nu\beta\beta$ decay rates upon new physics with sufficient precision, estimated to be 10–20%, to maximize the scientific reach of the planned experiments. Since $0\nu\beta\beta$ will be probed in a large nucleus, such as Xenon, these few-nucleon couplings to new physics have to be convolved with the background of nonperturbative nuclear physics. These calculations will require software development of architecture-aware, performance-optimized routines that exploit data- and task-parallel models optimized for extreme scales. This will require a tensor contraction engine capable of handling the exponentially more complex lattice calculations involving few-nucleon systems with mixed dependencies upon FFTs.

2. Computational and Data Strategies

2.1. Approach:

Lattice QCD is a numerical technique in which space-time is discretized, and the path-integral dictating the nonlinear quantum dynamics of the quark and gluon fields is evaluated by Monte Carlo techniques. Ensembles of field configurations with a range of lattice spacings and volumes are required at each selected set of light-quark masses in order to extrapolate to the continuum and infinite volume. Resources requirements increase with decreasing lattice spacing and increasing volume. Each configuration is saved to disk and archived for use in subsequent calculations. Large-scale configuration generation occurs over a period of months to years and requires capability HPC. A large number of light-quark propagators and low-lying eigenvectors are determined on each configuration through solution of large sparse linear systems. Algebraic multi-grid and deflation of low-mode eigenvectors are used for this task. On a single-field configuration, the intermediate storage can reach 20 to 35 TB in size, and performance I/O becomes crucial on the large problem sizes. Propagators typically require HPC capacity resources; however, propagators in large volumes and small lattice spacing can require capability resources. Correlation functions, produced from propagators and hadronic blocks (intermediate stages of calculations), are written to disk, archived, and moved off-site for subsequent analysis that generally can be performed on modest-capacity resources. Their generation requires HPC capacity facilities but can be performed equally well on capability resources.

2.2. Codes and Algorithms:

Typically, the most compute-intensive aspect of lattice QCD calculations is the inversion of large, sparse matrices. Efficient solver code has been developed using Algebraic Multi-Grid (AMG) methods. Specialized libraries have been developed for GPU (QUDA) and MIC (QPhiX) architectures. There are several open-source lattice QCD code bases developed by the USQCD Collaboration that either already use these solvers (Chroma, MILC) or will do so shortly (Qlua). These software packages are suited for current and upcoming HPC architectures supporting hybrid parallelism, vectorization, and memory optimizations.

For the Exascale era, significant improvements are needed in strong scaling of AMG methods, requiring the identification and exploitation of additional sources of task parallelism. Another target for improvement is the development of an efficient tensor-contraction engine with efficient support for

FFTs. The latter is specific to nuclear physics applications, where the calculation of contractions can be formulated as a distributed sparse-dense tensor contraction problem. These calculations could utilize sophisticated methods such as graph partitioning and dynamic load balancing. Furthermore, lattice QCD codes are designed for a sequential work flow, necessitating a growing I/O and data transit overhead between various steps. For Exascale computing, a significant redesign and refactoring of the current software stack are needed to support better modularity. New modules are needed to support new computational strategies and in-transit data analysis, as well as to employ scalable databases to streamline data analysis work flows and facilitate collaboration and data sharing.

3. Current and Future HPC Needs

3.1 Computational Hours:

In order to achieve sufficiently large physical volumes at small lattice spacings of approximately 0.045 fm, which are required for a proper continuum extrapolation, lattices with at least $L = 200$ are necessary. In the $0\nu\beta\beta$ decay calculations, only ensembles $m_\pi L > 6$ and spatial extent 8 fm will be used; calculations on the ones with lattice spacing 0.045 fm will be completed by 2020, and the remaining will be completed by 2025. In EDM calculations, where the small momentum limit is required, lattices with large volumes, 8 fm, will be used for analysis up in 2020, and up to 10 fm in 2025, while also studying the continuum limit down to lattice spacings of 0.045 fm in the former and 0.060 fm in the latter.

Most of the computational hours will be consumed in quark propagator calculation. Estimates are based on the current performance of the QUDA AMG on current generation architectures. A linear volume scaling and constant dependence on pion mass are assumed. Improvement factors x2 and x8 per GPU are anticipated by 2020 and 2025. Monte-Carlo statistics required for reliable signal is hard to estimate. We expect that between 10^5 and 10^6 samples will be needed. Based on the plans for gauge generation, every trajectory will be sampled as much as possible by shifting the origin of the correlation functions. We expect that $L^3/(2 m_\pi L)^3$ samples per time slice on the same configuration will be sufficiently independent. We anticipate sampling eight time slices on every gauge configuration.

3.2 Parallelism:

Currently, correlation function calculations mainly exploit data parallelism in contractions. Further opportunities for task-based parallelism will be investigated by fusing propagator generation and matrix element calculations together. Such fusion requires dynamic repartitioning of an MPI swarm depending on the current task, since different tasks may require different partition sizes to run efficiently. These two stages can be overlapped in heterogeneous architectures on the microscopic level (e.g., concurrency between GPU and CPU) and on the macroscopic level (e.g., concurrency between different machine partitions, optimized for different purposes).

The complex algebra and data dependencies in contractions will make manual control of subpartitioning and concurrency very tedious. We envisage development of a high-level domain-specific language, in which one will be able to express the algebra of propagator calculations and contractions. Analysis of data dependencies will provide necessary information to automatically isolate independent parts of the calculation and to partition available resources for execution of subtasks, while ensuring balanced load and concurrency. Combined with auto-tuning, such a framework will be a tremendous step forward in efficient utilization of Exascale architectures.

3.3 Memory:

For the matrix element calculations, the amount memory per node is not important. Both propagator and contraction calculations are embarrassingly parallel, and only a very small amount of communication is involved. In general, the total memory consumption increases linearly with the lattice volume and can be mitigated by running these calculations on large machine partitions. However, memory bandwidth is much more crucial as most of the calculation is spent in calculating ZGEMM operations.

3.4 Scratch Data and I/O:

Computing propagators and matrix elements is a capacity workload and embarrassingly parallel. The scratch space requirement is determined by the size of gauge configurations kept on disks at the same time, i.e., by the number of concurrent post-processing streams. It is realistic to run 64 and 128 parallel streams in parallel in 2020 and 2025 respectively. Including space required for output results, this amounts to about 50 TB and 350 TB scratch storage, respectively.

The work flow is not sensitive to I/O because I/O operations can be fully overlapped with computations. In the table we state 15–20% I/O time, which is typical when computation is not overlapped with I/O, and presented here only as the worst-case scenario. High-bandwidth I/O such as provided by burst buffers will still aid in applications where storing and reading eigenvectors may be required.

3.5 Long-term and Shared Online Data:

The data to be shared with possibilities for interactive data analysis depends on the project. For EDM and proton decay matrix elements, condensed and binned data may reach O(1 TB). For $0\nu\beta\beta$ decay, the data will include only five-point correlation functions, with approximately 10 GB of the online data storage required and thus negligible compared to other needs.

3.6 Archival Data Storage:

Although quark propagators will not be saved, the raw, unbinned correlation function information for EDM and proton decay matrix element calculations may reach between 10 TB and 100 TB. For $0\nu\beta\beta$ decay calculations, only the final, post-processed files will be archived, so the storage requirements are equal those discussed in Section 3.5.

3.7 Workflows:

As described above, lattice QCD nuclear physics calculations can be divided into three steps: ensemble generation, measurements, and data analysis. The first two require massive amounts of computational resources, whereas the last is a data management/bookkeeping problem. Generating the Markov chain of ensembles is a sequential process and can only be accelerated by running on as many nodes as strong scaling allows. The resulting configurations will be kept to be re-used in future calculations.

Step two consists of two parts: (i) computing quark propagators and (ii) contracting them together into nucleon correlation functions containing relevant nucleon matrix elements. At the moment, computed quark propagators are temporarily stored to a scratch disk, and read by another job that performs contractions for NN matrix element. Contractions for single-nucleon correlation functions (EDM, proton decay) are simpler and are performed on the fly. In nuclear physics calculations at the physical pion mass, approximately 10^6 quark propagators have to be generated and contracted.

The third step (data analysis) can be improved by employing data management systems, such as scalable, in-memory databases. They would allow fast exploration of analysis techniques and simplify collaboration and data sharing.

3.8 Many-Core and/or GPU Readiness:

Solver packages used in HMC and quark propagator calculation, such as QPhiX and QUDA, are well optimized for MIC and GPU use, respectively. The FFTs currently use MKL and cuFFT, and these libraries are well supported and optimized for the targeted architectures. We expect that future versions of cuFFT will include nvlink support and algorithmic strong scaling improvements. Our contraction routines utilize batched cuBLAS and MKL-BLAS routines under the hood, and more than 80% of the time is spent in those routines. Therefore, we consider most of our code well ready for Exascale nuclear physics applications.

Contractions that are needed for nEDM, proton decay, and neutron–antineutron oscillations are implemented, depending on software package, either in QDP++, which may be run on GPUs within JIT framework, or in QDP/C, which is OpenMP-ready and can run efficiently on MIC.

3.9 Software Applications, Libraries, and Tools:

Lattice QCD software stack (qmp/qio/qdp++/Chroma) requires C++11, better C++14/17 for compilation of parts of the software; LLVM/Clang for code optimizations and JIT compilation within the QDP-JIT framework; and nvcc and the CUDA SDK for CUDA capable code. More JIT features in upcoming compiler infrastructures, as well as better tools for debugging JIT-enabled code (for example, attaching debug symbols) are highly desirable. Allinea DDT, TotalView, vTune, and Vampir are expected to be available for debugging and profiling.

Efficient implementation and support of the MPI-standard are essential, with improvements for asynchronous nearest neighbor communications and specialized routines to improve concurrency and latency. We also hope for thread-parallel performance improvements in terms of re-using existing threads inside these libraries from outside applications to reduce thread fork and merge overhead.

We will rely on the following architecturally optimized libraries: BLAS+LAPACK (cuBLAS, MKL) for small linear algebra tasks; distributed FFT libraries for accelerating some tensor contractions. For parallel I/O, our codes employ parallel HDF5. It is desirable to control the topological awareness of the underlying MPI-IO part, such as respecting distances to Lustre OSTs.

3.10 HPC Services:

Nonstandard services we will need for this project are scalable databases that can be tied to a web interface for collaboration and data-sharing purposes. Training in code optimization for upcoming, especially for heterogeneous, architectures will be advantageous. Code-versioning systems with support for continuous integration, unit testing and coding style standardization, and versioning systems with fine-grained access controls will improve development efficiency.

3.11 Additional Needs:

Lattice calculations are often distributed among several (LCF and local) computational centers. In the Exascale era, local computer clusters will be powerful enough to host parts of the post-processing work flows and data analysis. Tighter integration between these will improve productivity. One can think about scheduling data transfers, analysis, and post-processing based on data availability. For example, as soon as a configuration is produced at center A, it is transferred to center B to perform the

measurements and data analysis, streaming intermediate results into a database at center C. A database that integrates all repository allocations at the various centers and national laboratories, and a unified job submission system would improve transparency for the users, as well as the committees deciding CPU time allocations.

Requirements Summary Worksheet

1. $0\nu\beta\beta$ -decay

Code: _____	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)*	60M h	18M h (0.3x)	1.2M h (0.02x)
Computational node hours (homogeneous many-core)**	-	-	10M h
Computational node hours with (w/GPU or accelerator)***	-	32M h	64M h
Memory per node	8 GB	24 GB (3x)	32 GB (4x)
Aggregate memory	0.1 TB	5 TB (50x)	20 TB (200x)
Data read and written per run	1 TB	15 TB (15x)	30 TB (30x)
Maximum I/O bandwidth needed	5 GB/sec	5 GB/sec (1x)	5GB/sec (1x)
Percentage of runtime for I/O	15–20%	(0.1x)	(0.1x)
Scratch file system space needed	2 TB	50 TB (25x)	200 TB (100x)
Permanent online data storage	0.02 TB	0.04 TB (2x)	0.04 TB (2x)
Archival data storage needed	0.02 TB	0.04 TB (2x)	0.04 TB (2x)

* Core hours are used for conventional processors. (i.e., node-hours × cores per node). Intel Ivy Bridge is an example.

** Node hours are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi Knights Landing is an example.

*** Node hours are used for GPU or accelerator usage.

**** For example, 32 times column 1.

2. Neutron EDM and Proton Decay

Code: _____	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)*	Can use instead of GPUs	Can use instead of GPUs	Can use instead of GPUs
Computational node hours (homogeneous many-core)**	Can use instead of GPUs	Can use instead of GPUs	Can use instead of GPUs
Computational node hours (with GPU or accelerator)***	2M h	87.4M h (44x)	510M h (255x)
Memory per node (per GPU)	8 GB	Scale as total GPU memory (1x)	Scale as total GPU memory (1x)
Aggregate memory	2 TB	10 TB (5x)	100 TB (50x)
Data read and written per run	0.01 TB	0.015 TB (1.5x)	0.075 TB (7.5x)
Maximum I/O bandwidth needed	5 GB/sec	1 X	1 X
Percentage of runtime for I/O	Negligible	Negligible	Negligible
Scratch file system space needed	Not required	Not required	Not required
Permanent online data storage	0.05 TB	0.06 TB (1.2x)	0.11 TB (2.2x)
Archival data storage needed	10 TB	22 TB (2.2x)	90 TB (9x)

*“Core hours” are used for “conventional” processors. (i.e., node-hours × cores per node). Intel “Ivy Bridge” is an example.

**“Node hours” are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi “Knights Landing” is an example.

***“Node hours” are used for “GPU or accelerator” usage

**** For example, 32 times column 1.

The Gluonic Structure of Matter

Lead Authors: David Richards (Thomas Jefferson National Accelerator Facility), Martin Savage (Institute for Nuclear Theory)

Contributors: Robert Edwards (Thomas Jefferson National Accelerator Facility), Balint Joo (Thomas Jefferson National Accelerator Facility) and Kostas Orginos (College of William and Mary and Thomas Jefferson National Accelerator Facility)

1. Description of Research

1.1 Overview and Context:

Gluons are the force carriers of Quantum Chromodynamics (QCD) and play a pivotal role in generating the mass and structure of the proton and of nuclei. Precision studies of their role are emerging as the next frontier in hadronic physics. The numerical technique of lattice QCD, in which QCD is systematically solved in a finite to and discretized volume of space-time, provides the means to investigate the ways that gluons manifest themselves. The questions that lattice QCD can address are profound: What is the spatial distribution of gluons within a nucleon? How does the orbital motion of gluons and sea quarks give rise to the spin of the nucleon? How do gluons contribute to the mass of the proton and other hadrons? What does the gluonic structure of hybrid mesons and baryons reveal of the role of gluons? Is the distribution of gluons in a nucleon bound in a nucleus different from that in a free nucleon? In an analogous way to how lattice calculations of quark structure have advanced the experimental hadron structure programs at CEBAF and RHIC-spin, so these will capitalize on the exploration of the role of gluons in the GlueX experiment at JLab@12GeV and at a proposed Electron-Ion Collider (EIC).

1.2 Research Objectives for the Next Decade:

The gluonic contribution to the confined energy that constitutes the mass of the proton, nuclei, and other particles will be precisely determined, complementing the extensive body of analogous work on contributions from the quarks. It will be further extended to include the structure of excited particles called hybrids in which the gluons manifest themselves as explicit degrees of freedom. The x -dependence of the unpolarized-parton distribution functions associated with the proton and nuclei will be determined in order to reveal the contribution of the gluonic degrees of freedom at small Bjorken- x . The contribution of quark and gluon spin and of orbital angular momentum to the spin of the nucleon is encoded in the three-dimensional Generalized Parton Distributions (GPDs) and Transverse-Momentum-Dependent distributions (TMDs). We will perform direct calculations of the x -dependence of the GPDs and compute the TMDs to reveal the contributions of the gluons to the spin of the nucleon. Each of these studies will be performed with continuum and infinite-volume extrapolations to provide quantifiable and controlled uncertainties.




Figure 2. The spin-decomposition of the proton. [Permission to reproduce this figure given by Keh-Fei Liu.]

2. Computational and Data Strategies

2.1 Approach:

Ensembles of gluon configurations with a range of lattice spacings and volumes will be produced to study the gluonic structure of matter, which require capability computing resources. They will have lattice spacings small enough to enable the study of high momentum correlation functions and spatial and temporal extents sufficient to fully quantify finite-volume and thermal systematic effects.

Electromagnetism is expected to make only small contributions to the gluonic quantities of importance and need not be included. Hadronic correlation functions on these ensembles are required for determining gluonic structure and are calculated independently on each configuration. Quark propagators representing the parallel transport of quarks on the configuration are required; typically, many tens of thousands of these will be computed on each configuration, and these can be computed with high efficiency using AMG methods, which are saved for subsequent use. Two- and three-point correlation functions are calculated through tensor contractions that are stored for subsequent analysis. Correlation functions typically require capacity computing resources, but for particularly large configurations, capability resources are becoming necessary.

Two- and three-point correlation functions are calculated through tensor contractions that are stored for subsequent analysis. Correlation functions typically require capacity computing resources, but for particularly large configurations, capability resources are becoming necessary.

2.2 Codes and Algorithms:

Discretization of the quark fields generates nearest neighbor stencil-like terms that are sensitive to communications performance, and these algorithms face scaling issues in fat-node systems, such as GPU-based architectures. Methods using deflation and AMG preconditioning tame the condition number and are used to calculate quark propagators. The overlap discretization is potentially highly demanding, but admits mass shifting whereby several valence quark masses can be computed for




Figure 3. Gluons constitute much of the energy confined into a proton; the question is, How much?

essentially the cost of the lightest, and has far simpler operator mixing for matrix elements. Correlation functions require a series of tensor contractions, which can be performed independently for each configuration and time slice.

The **Chroma** lattice field theory suite was designed as part of USQCD's SciDAC initiative and has evolved to meet physics goals and computer architectures. It is written in terms of the QDP++ data parallel layer and is re-implemented in the QDP-JIT package, supporting hybrid OpenMP-MPI parallelism and MPI+PTX for GPUs and now the LLVM framework. High-performance libraries implement solvers that are tuned to specific architectures, such as QUDA for GPUs, QPhiX for Xeon & Xeon Phi, and QOP-MG for all-MPI builds on CPUs. There is a three-dimensional analogue called **harom** that facilitates the calculation of time-sliced correlation functions, and **Redstar** supports many-body operator constructions and automates the optimal edge reduction of graphs. It is built over an API that supports threaded tensor contractions optimized for CPUs, GPUs, and KNLs. The overlap fermion program uses a bespoke code base that has been implemented on GPUs in addition to CPUs.

3. Current and Future HPC Needs

3.1 Computational Hours:

Current projects are calculating the contribution of valence and sea quarks to the structure of the nucleon using isotropic clover action and the gluon contribution using overlap fermions on domain-wall-fermion lattices. Gauge generation is, in general, performed on accelerated nodes, with around 1.9M GPU hours annually dedicated to clover-gauge generation on Titan. The current projects consume around 200M core hours annually on leadership-class and USQCD facilities. We use as an exemplar the Wilson-clover calculation for future needs. In 2020, the physics goals are the mass decomposition of the pion and nucleon through the calculation of the stress-energy tensor and the direct calculation of the x-dependence of GPDs and of TMDs, with controlled uncertainties due to the finite volume and finite lattice spacing. The generation of one ensemble (spatial volume of 6 fm, lattice spacing of 0.06 fm) will require 101M GPU hours, and the analysis phase, comprising the calculation of propagators and contractions, will require 160M GPU hours. In 2025, the physics goals will comprise the mass composition of excited states and, in particular, hybrids and the calculation of the modification of gluon structure in nuclei. Both of these projects require the generation of lattices at larger volumes. The cost of one ensemble (spatial volume of 8 fm, lattice spacing of 0.06 fm) will be 268M GPU hours, with the analysis phase requiring 900M GPU hours.

3.2 Parallelism:

The analysis phase uses throughput computation of the Dirac equation solution vectors for many right-hand sides on numerous, but small, partitions of the system, which are subsequently partially contracted to form tensors for each time slice of the lattice, and must either be paged to disk or held in memory. The contraction step, which consumes the tensors in memory or on disk, typically uses single many-core or GPU nodes for each time slice of the contractions. The Hadron node code relies on batching to optimize the tensor contractions. Here, each spin component of the tensor contraction is handed to a smaller part of the node, such as a numa node, to increase memory and performance within that smaller partition.

3.3 Memory:

The Dirac solvers are run in small partitions to increase their efficiency, limited only by memory size per node; we currently use GPU and KNL nodes with up to 128 GB of main memory. The solver codes use specialized caching routines to page data in/out of fast memory to slower main memory, and we expect this to remain essential. The contraction step is also memory limited and requires the most care in the work flow. The tensors occupy approximately 11 GB on a single computational node for the anticipated 6 fm spatial sizes. Future developments will switch to a sparse (banded) structure for which we expect about 1 GB will be needed. The code will have to page these to intermediate storage.

3.4 Scratch Data and I/O:

There are a large number of solution vectors that must be processed and saved on disk in a form, “perambulators,” that allows them to be re-used in subsequent contractions. For the 2020 lattices, they are 1 TB per configuration. The computation of matrix elements requires the solution vectors be written in the propagator phase and read back to form “generalized perambulators.” These latter require about 150 TB per configuration for the 2020 calculation and 900 TB for the 2025 calculation. While the tensor contractions are executed, there will be intermediate temporaries created and re-used, which will reduce the overall computational cost. Typically, a few hundred of these temporaries are anticipated, about 1 GB each, to be created per time slice (computational node). The I/O bandwidth required to write and read these objects is expected to be large enough that some kind of rapid paging to burst buffers or network database will be needed. The trade-off for the computations is the paging of these temporaries versus recomputing them to minimize the overall wall-clock time.

3.5 Long-term and Shared Online Data:

The long-term storage for the perambulators and generalized propagators will be shared with other projects in nuclear physics.

3.6 Archival Data Storage:

All perambulator and generalized perambulator data across all ensembles in the 2020 and 2025 gluonic campaigns will be stored.

3.7 Workflows:

The calculation of the perambulators, and the generalized perambulators needed for matrix element calculations, is on GPU nodes. The contraction to form correlation functions is performed on CPUs as a separate job. We envision both steps being performed on GPUs or many-core CPUs, possibly as part of a single job, in 2020 and 2025.

3.8 Many-Core and/or GPU Readiness:

Extensive software development has been under way to support many-core systems and GPU systems. The Dirac solvers is one area of active development under SciDAC and in collaboration with Nvidia and Intel. We are actively extending our basic data-parallel software framework used by Chroma and QDP++ as a part of SciDAC. The other area is support for tensor contractions within our Hadron package. We expect to take advantage of these systems in the future.

3.9 Software Applications, Libraries, and Tools:

We rely heavily on the software stack developed under SciDAC and seek continued collaborations with industry and also laboratory partners for continued innovation. There is a continued need for BLAS routines, such as ZGEMM, optimized for future architectures. C++ support is required.

3.10 HPC Services:

Possibilities might include advanced training, data analytics and visualization assistance, support servers, collaboration tools, web interfaces, federated authentication services and gateways.

3.11 Additional Needs:

Training provided for graduate students and postdocs at our host institutions, for example, as a part of SciDAC, is crucial to develop software, and exploit HPC facilities. Local computing resources have been a critical part of the success of USQCD. We have relied heavily on the leadership resources for our most demanding computations, and the local compute facilities for the continued analysis of this data. We rely on this mutually beneficial coordination to continue into the future.

Requirements Summary Worksheet

Code: Chroma and Redstar	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)****	Future Usage: 2025 (As a factor of column 1)****
Computational core hours (conventional)*	200M		
Computational node hours (homogeneous many-core)**	2M	Can use either CPU or GPU	Can use either CPU or GPU
Computational node hours (with vGPU or accelerator)***	2M	Gauge: 101M (x50) Prop: 160M (x80)	Gauge: 268M (x134) Prop: 900M (x450)
Memory per node	128 GB	128 GB	128 GB
Aggregate memory	8 TB	100 TB	300 TB
Data read and written per run	1 TB	100 TB	300 TB
Maximum I/O bandwidth needed	GB/sec	GB/sec	GB/sec
Percentage of runtime for I/O	10	20	20
Scratch file system space needed	10 TB	200 TB	900 TB
Permanent online data storage	100 TB	150 PB	900 PB
Archival data storage needed	500 TB	1,200 PB	8,800 PB

* Core hours are used for conventional processors. (i.e., node hours × cores per node). Intel Ivy Bridge is an example.

** Node hours are used for homogenous many-core architectures. A self-hosted Intel Xeon Phi Knights Landing is an example.

*** Node hours are used for GPU or accelerator usage.

**** For example, 32 times column 1.

Configuration Generation

Lead Authors: Balint Joo (Thomas Jefferson National Accelerator Facility), Thorsten Kurth (Lawrence Berkeley National Laboratory and National Energy Research Scientific Computing Center), Robert Edwards (Thomas Jefferson National Accelerator Facility)

Contributors: William Detmold (Massachusetts Institute of Technology) and Kostas Orginos (College of William and Mary and Thomas Jefferson National Accelerator Facility)

1. Description of Research

Lattice QCD calculations impacting Nuclear Physics proceed by computing observables on lattice gauge configurations, which are snapshots of the QCD gluonic field in the vacuum (akin to the electric and magnetic fields in electromagnetism). QCD configurations include all gluon self interactions and sea-quark gluon interactions. Once ensembles of configurations are generated, the remainder of a calculation, referred to as analysis or measurement, involves the computation of quark propagators (the QCD analogue of placing a test charge in an electromagnetic field) and combining them to form the necessary observables. Measurement strategies are different for the various calculations and are described in detail in other case studies. In this case study, we focus on gauge generation, which is common for many calculations.

2. Computational and Data Strategies

2.1 Approach:

In order to allow precision calculations in nuclear physics, the ensembles need to satisfy certain properties: i) the quark masses in the calculation must match their physical values; ii) the physical lattice volumes need to be sufficiently large to avoid unwanted size and thermal effects; iii) at least 3 lattice spacings are needed for controlled continuum extrapolation; and iv) the effects of isospin breaking and electromagnetic interactions need to be included for certain observables.

Current calculations cannot simultaneously satisfy all the requirements above. Calculations for cold nuclear physics in the US are typically carried out with degenerate light quarks (ignoring isospin breaking), without the inclusion of electromagnetic effects and using unphysically heavy light quarks, characterized by π -meson mass of approximately 200 MeV, rather than the physical value $m_{\pi}^{\text{Phys}} \sim 135$ MeV due to the high computational cost of meeting these requirements. Only recently have isospin breaking and electromagnetism been introduced into generation algorithms, but due to the costs introduced such calculations have been limited to relatively small volumes when using light quarks ($m_{\pi}L \sim 4$, $m_{\pi} \sim 197$ MeV). An additional challenge expected in the next 10 years is a growth in the cost of generating independent configurations as the lattice spacing is reduced, making ensembles with finer resolution computationally much more expensive than current calculations. The data needs of the gauge generation are relatively modest, compared to the computational cost.

2.2 Codes and Algorithms:

The key algorithm for gauge generation is the hybrid Monte Carlo algorithm (and its variants). This is a Hybrid Molecular Dynamics Markov Chain method, which proposes new gauge configurations from previous ones by utilizing Hamiltonian Molecular Dynamics (MD) in a fictitious MD-time. The new configurations are subjected to Metropolis Accept/Reject steps. Area-preserving and time-reversible MD integrators are used in order to satisfy detailed balance, which along with ergodicity is sufficient for the algorithm to equilibrate to and sample the correct equilibrium probability distributions. At each step of the MD, the forces due to the terms in the action need to be evaluated. For sea quarks, these involve the solution of the lattice Dirac equation governing quark–gluon interactions. The solution of the Dirac equation is the dominant cost of current calculations, consuming ~90% of the time in the gauge generation program.

The solvers used in HMC have traditionally been simple Krylov subspace iterative methods for sparse systems such as Conjugate Gradients (CG) or Stabilized BiConjugate Gradients (BiCGStab). The cost of these methods is driven by the condition number of the lattice Dirac operator which diverges as $1/(am_q)$, with m_q being the quark mass. Due to the nearest neighbor stencil like components of the Wilson-Clover formulation of quarks; which are very sensitive to communications performance; the CG and BiCGStab algorithms have faced strong scaling issues in recent fat-node systems such as GPU based architectures. Recent developments have attempted to address both these issues: Solvers with Domain Decomposed (DD) preconditioners have demonstrated superior scaling performance to CG and BiCGStab. Methods using Deflation and Adaptive Multi-Grid (AMG) preconditioning have been developed to deal with the problem of the divergence of the condition number, although these have so far found best use in measurement where their setup costs could be amortized over very many solves. The addition of AMG to gauge generation is a recent development that promises to make the cost of calculations essentially independent of the quark masses.

The main workhorse code suite for Gauge Generation in cold Nuclear Physics calculations in the US is the Chroma software suite developed by the USQCD collaboration through SciDAC funding. Chroma is written in terms of a data-parallel, domain-specific framework called QDP++ that has implementations (QDP-JIT) enabling Chroma to run either on conventional or many-core CPU systems as well as GPU-based systems. On node parallelism is expressed with OpenMP for CPUs and CUDA or PTX for GPUs. Additionally, Chroma has been integrated with high-performance, domain-specific solver libraries to further increase code efficiency on specific hardware architectures, such as QUDA for GPUs, QPhiX for Xeon and Xeon Phi and QOP-MG for all-MPI builds on CPUs.

3. Current and Future HPC Needs

3.1 Computational Hours:

Gauge generation currently takes place at the Leadership Computing facilities utilizing INCITE allocations and occasionally ALCC hours. The majority of nuclear physics allocations have been on Titan at ORNL, and for this reason we will base our reasoning on GPU-based systems, though we emphasize that we are agnostic about architectures. The portion of the current INCITE allocation devoted to gauge generation for cold Nuclear Physics is 58M Titan core hours. In addition approximately 120M Titan core hours have been used over the last year from two ALCC projects, giving a total current usage 178M Titan core hours, or approximately 6M GPU hours.

The costs of generating 1,000 independent electromagnetically neutral configurations are shown in Figure 1. Calculations with a 6-fm spatial volume are useful for spectroscopy and hadron structure calculations, whereas larger 10-fm spatial volumes will be needed for determinations of the Nuclear Forces. The costs were extrapolated from existing production calculations on Titan and assume an additional cost reduction factor of two based on near term “line of sight” algorithmic improvements. The cost of including electromagnetism and isospin breaking is estimated to be an additional factor of 10 times.




Figure 1: Cost in M-GPU hours of generating electromagnetically neutral configurations at three physical volumes, with temporal extent of approximately 17 fm. The requirements for including electromagnetism and isospin breaking are assumed to be approximately 10 times higher.

For concreteness of discussion, we consider, in this case study, a scenario where in 2020 1,000 independent configurations of electromagnetically neutral lattices are produced with spatial extents of 6 fm, 8 fm, and 10 fm at lattice spacings of 0.077 fm (the right-most points on the lines in Figure 1) and 0.06 fm. Between 2020 and 2025, we anticipate that ensembles corresponding to every point in the diagram will have been generated and production of electromagnetically charged ensembles will have begun, with the same spatial extents at the larger lattice spacing. In 2025 we will embark on generating electromagnetically charged configurations at the finer lattice spacing, and we will be able to finish generating the smaller two ensembles (1,000 independent configurations in each), and we will be about 10% of the way through generating the larger (100 configurations). To achieve this would require 1,125 M Titan GPU hour equivalents in 2020, and 4,300 M Titan GPU hour equivalents in 2025.

3.2 Parallelism:

The lattice structure in lattice QCD, combined with locality, provides a large degree of data parallelism. Typical lattices volumes in the 2020–2025 time frame are expected to have between 128 and 256 spatial sites, giving approximately 10^{12} -way data parallelism. This is exploited by domain decomposition of the lattice among MPI tasks and threading or accelerating the work on the local domain in each MPI task. To overcome strong scaling challenges going to exascale systems, additional parallelism may need to be exploited, for example, among stencil directions. Strong scaling is expected to be the greatest challenge, especially for the AMG solvers.

3.3 Memory:

Currently our Gauge Generation in systems with spatial extent 4.9 fm fits comfortably on 1,024 nodes of Titan. In 2020 we plan to work with lattices that are 12 times the size of the volume we utilize today, whereas in 2025 we would require lattices 82 times the size used today. In between these years, however, we would need to generate lattices that are up to 33 times the size of current lattices, and so we will take this as the requirement for 2020. Bearing this in mind, in 2020 we would need an upper bound on total memory for our calculations of 1,056 TB DDR and 198 TB fast memory, which could be satisfied with approximately 2,200 nodes of a Summit-like system with 512 GB DDR per node and 96 GB fast memory. In 2025 we will need 2,624 TB DDR and 492 TB fast memory, which could be met with ~2,600 nodes, each with 1,024 GB DDR and 192 GB fast memory. We emphasize that relatively small partitions of high-memory nodes, such as dense GPU nodes, are not a requirement, and this need could also be met with larger partitions of less-memory-rich nodes.

3.4 Scratch Data and I/O:

Current calculations spend between 0.5% and 2.5% of their run time in I/O writing out configurations. A typical configuration of $64^3 \times 128$ sites takes 19 GB on disk. Scaling the lattice volumes and considering the approximate speed of simulation enabled by partitions of the size discussed in Section 3.3, we estimate requiring a disk read/write bandwidth of about 25 GB/sec. In 2020, a 24-hour calculation is expected to output about 20 TB of data, and scratch space for 3 streams (60 TB) should be adequate. In 2025, the data written will be larger per configuration; however, due to the increased lattice sizes, the calculations will proceed more slowly and overall a 24-hour run will still be likely to produce 20 TB and a 60 TB scratch space and, as in 2020, will be adequate.

3.5 Long-term and Shared Online Data:

We intend to make the generated gauge configurations available online in the long term for measurement. Typically this would not need to include all the trajectories but only a subset of them (typically every 10), although the separation can grow as the autocorrelation time grows. Our current runs when they get to 500 configurations would constitute a data set of ~30 TB. Our working scenario envisages generating 1,000 configurations each in the 6-fm, 8-fm, and 10-fm ensembles in 2020 with lattice spacings of 0.077 fm and 0.06 fm, giving a cumulative long-term storage need of 1.7 PB. By the end of 2025, including the configurations we intend to generate in the intervening years from 2000, the cumulative online storage requirement for the gauge configurations would grow to 7.8 PB.

3.6 Archival Data Storage:

While we only need to make the "independent" configurations available online, it is useful to archive every saved trajectory, including the checkpoints. This can be helpful if a simulation needs to be restarted or for autocorrelation analysis. Typically, every two trajectories are checkpointed. With this in mind and taking into account the number of trajectories needed for the various ensembles, we anticipate an archival need of around 150 TB for current-day simulations, approximately 12.6 PB for 2020, and approximately 85.2 PB for 2025.

3.7 Workflows:

The gauge-generation work flow is straightforward. Configurations are generated in sequence, checkpointing every two trajectories. The checkpoints are the same as the main generated data and can be moved to either online storage or archival once written. Additionally it is possible to carry out several independent streams of gauge generation concurrently (e.g., at different parameters or having started from independent configurations). In the 2020 time frame, a modification we can envisage is that

configurations may be passed to coscheduled jobs, e.g., via burst buffer or system fabric for I/O to online storage and rudimentary initial analysis.

3.8 Many-Core and/or GPU Readiness:

Our software stack for HMC is GPU ready, through a combination of Chroma, built over QDP-JIT combined with the QUDA library. Multicore readiness is proceeding via Chroma, built with QDP-JIT/LLVM and the QPhiX solver library.

In general, however, our software frameworks have been designed prior to the era of many-core and accelerator based systems with long SIMD lengths. Going forward to the Exascale, a significant software effort needs to be invested in the creation of new software frameworks with a strong emphasis on performance portability and productivity. We are engaged in these efforts through Application Readiness partnerships (NESAP, Theta ESP) and share our experiences at the community level by participation in meetings such as HPCOR'15 and the DOE Centers of Excellence workshops on Performance Portability.

3.9 Software Applications, Libraries, and Tools:

Certain libraries are needed such as BLAS and LAPACK. We anticipate using either C/C++ or languages that compile into C/C++ in the 2020–2025 time frame. Good support for emerging standards, such as evolutions of OpenMP, OpenACC and the C++ standard (C++11/14/17 and beyond), will be needed. Our current software approach also requires the LLVM compiler framework library to enable the just-in-time compilation of QDP-JIT for non-PTX architectures.

3.10 HPC Services:

With the increase in the amount of data being generated and to enable multi-center campaigns, we will need to be more formal about data management, including setting up curated online libraries of ensembles of gauge configurations and tools to enable straightforward transfer of data between sites (e.g., evolutions of Globus/GlobusOnline) and sharing among measurement projects. It would also be useful if this were synergistic with existing configuration libraries (e.g., the gauge connection at NERSC and the ILDG, for data that have been publicly published).

3.11 Additional Needs:

In order for these projects to be successful, getting an early, well-informed start to application porting and readiness activities for upcoming systems is and will remain critical. We can point to the success of the NERSC NESAP program for helping us prepare for Cori and KNL systems, and activities such as OLCF CAAR and ALCF ESP also aim to fill exactly this need. We need more of these kinds of activities, as well as early access to technical information, expertise, early-user experiences, and best practices, which can be disseminated through webinars and workshops. Development work requires access to emulators/simulators, software tools, and prototype systems.

Requirements Summary Worksheet

Code: Chroma	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	We can use instead of GPUs	We can use instead of GPUs	We can use instead of GPUs
Computational node hours (homogeneous many-core)	Can use instead of GPU	Can use Instead of GPU	Can use Instead of GPU
Computational node hours (with GPU or accelerator)	6M	~187x (= 1,125M GPU h)	~720x (= 4,300M GPU h)
Memory per node	32 GB DDR + 6 GB fast	16x (512 GB DDR + 96 GB fast)	32x (1,024 GB DDR + 192 GB fast)
Aggregate memory	32 TB DDR +6 TBG fast TB	33x (1,056 TB DDR + 198 TB fast)	82x (2,624 TB DDR + 492 TB fast)
Data read and written per run	0.9 TB	20x (20 TB)	20x (20 TB)
Maximum I/O bandwidth needed	1 GB/sec	25x (25 GB/sec)	25x (25 GB/sec)
Percentage of runtime for I/O	1–2%	1x (1%)	1x (1%)
Scratch file system space needed	8 TB	~7.5x (60 TB)	7.5x (60 TB)
Permanent online data storage	30 TB	~57x (1,700 TB)	~237x (7,800 TB)
Archival data storage needed	150 TB	~84x (12,600 TB)	~568x (85,200 TB)

D.5 Case Studies Addressing Hot Quantum Chromodynamics

Quark Production and Equilibration Processes in Non-equilibrium QCD

Lead Authors: Sören Schlichting and Björn Schenke (Brookhaven National Laboratory)

1. Description of Research

1.1 Overview and Context:

Over the past decades, experiments at the Relativistic Heavy Ion Collider (RHIC) have revealed exciting properties of nuclear matter under extreme conditions, including the existence of a new state of matter, the Quark-Gluon Plasma (QGP). However, the theoretical understanding of these observations remains a challenge and understanding how the Quark-Gluon plasma is formed from the earliest moments in a nuclear collision has been identified as one of the key questions in the 2015 Nuclear Science Long Range Plan.

The central objective of this project is to develop an ab-initio description of high-energy heavy-ion collisions, which establishes a first-principles connection between collider experiments and the underlying theory of quantum chromodynamics (QCD). So far the focus of this line of research has been to understand different aspects of the initial state, the early-time dynamics, and the thermalization process in high-energy QCD by use of real-time lattice gauge theory techniques. While present simulations have been able to capture the dynamics on a qualitative level, further progress in theory and computing is required to resolve this long-standing question and obtain quantitative insight into experimentally relevant questions.

1.2 Research Objectives for the Next Decade:

Over the next decade, significant improvements over the present state-of-the art can be anticipated based on the inclusion of dynamical quarks into existing simulations of non-Abelian gauge fields. Outstanding scientific questions, including a first-principles understanding of the mechanism of quark production, the electromagnetic response of the pre-equilibrium plasma, as well as the real-time dynamics of anomalous transport phenomena such as the chiral magnetic effect, can be addressed for the first time within this framework. Computationally, the inclusion of dynamical fermions marks a milestone, as the computational resources required for each simulation increases by a factor of typically 10^3 – 10^4 compared to pure gauge theory, leading to a significant increase in computational demands over the next decade.

2. Computational and Data Strategies

2.1 Approach:

Evolution equations for non-Abelian gauge fields are discretized on a spatial lattice and solved numerically. Dynamical fermions are included by expanding the fermion field in a complete operator basis at the initial time and computing the time-dependent wave functions by solving the discretized

evolution equations for each of the fermion modes. Physical observables are computed at different times of the evolution based on the solutions of the equations of motion for the gauge fields and fermion wave functions. Discretization effects are estimated by varying the lattice size and spacing.

2.2 Codes and Algorithms:

Software for non-Abelian gauge-field evolution with dynamical fermions has recently been developed, and first applications are starting to be explored. Codes are written in the C++ language and employ hybrid OpenMP/MPI parallelization. The work flow of the present software is as follows:

1. Initialize gauge-field configurations and wave functions for fermions.
2. Compute the gauge fields at the next time step via a leap-frog update (serial).
3. Compute fermion wave functions at the next time step, using leap-frog update (massively parallel).
4. Compute the fermion currents entering the gauge-field equations of motion (massively parallel with collective communication).
5. Iterate steps 2, 3 and 4 to obtain the time evolution.
6. Compute observables at times of interest (massively parallel with collective communication).

Computation of observables and initialization routines make use of the GSL and FFTW libraries; statistical averaging over initial conditions and data analysis are typically performed offline.

Currently, computation of $\gamma^\mu D_\mu \phi$ in the update of the fermion wave functions (step 3) requires the largest fraction of computational resources associated with these calculations, consuming approximately 80% of the wall time. Most of the residual computation time is spent in the computation of fermion currents (step 4) and fermion observables (step 5). Performance of the code is memory-bound, as nonlocal memory access is required in all steps of the calculation.

3. Current and Future HPC Needs

3.1 Computational Hours:

Non-equilibrium simulations without dynamical fermions are presently performed on sufficiently large lattices to extract continuum results ($\sim 128^3$) and consume 2M to 5M core hours per year. Simulations with dynamical fermions have only been developed recently and presently require a similar amount of 2M to 5M core hours computing time per year. However, since the computational requirements of a single simulation with dynamical fermions on a $24 \times 24 \times 64$ is ~ 150 k core hours, these studies have been limited to small lattices, where only drastically simplified situations involving either homogeneous or slowly varying gauge-field configurations can be studied.

Scientific goals require the present studies with dynamical fermions to be extended to significantly larger lattices ($\sim 128^3$) to ensure that all relevant scales of a realistic gauge-field configuration can be resolved. While the computational cost for the pure gauge-field evolution scales linearly with the size of the lattice $N_x \times N_y \times N_z$, the cost for including dynamical fermions scales quadratically in the lattice size as the evolution equations have to be solved at a complexity $N_x \times N_y \times N_z$ for each of the $N_x \times N_y \times N_z$ lattice sites. Based on this scaling, we project an increase in the computational cost of a single simulation by at least two to three orders of magnitude. Since for realistic field configurations statistical averages have to be computed with an ensemble of at least 10-50 different initial conditions. Depending on the observables of interest this will further increase the computational cost. Moreover, ensuring

convergence to continuum results, typically requires ~ 10 independent runs where the discretization parameters are varied. The cost of computations also depends on the fermion formulation. Performing calculations with chiral fermions, which is needed for the study of chiral magnetic effect, is at least 10 times more expensive than performing calculations with Wilson fermions. As an example, we estimate that the cost of running simulations with Wilson fermions on 96^3 lattices would take about 30M core hours. If we perform calculations for 50 different initial conditions, 1,500M core hours will be needed. Overall, we expect a rapid increase of the computational demand by more than 2 orders of magnitude through 2020 and a milder increase of about another order of magnitude through 2025.

3.2 Parallelism:

Since the evolution of individual fermion modes is independent, the problem exhibits a high degree of parallelism. Our current software uses a hybrid MPI and Open-MP parallelization concept, which for the present state-of-the-art simulations, runs on 1,536 nodes employing 24 physical processors per node plus hyper-threading. All future studies on larger lattices will require an increase in the number of computing nodes proportional to the increase in core hours, i.e. by two to three orders of magnitude.

3.3 Memory:

In present simulations, the aggregate memory consumption is ~ 20 TB, equally distributed among all available compute nodes. Our present production runs on 1,536 nodes require approximately 16 GB of memory per node, which could be easily decreased by increasing the total number of nodes without a significant impact on performance. All future studies on larger lattices will require an increase in aggregate memory consumption proportional to the increase in core hours, i.e. by two to three orders of magnitude. Since the performance of the code is limited by memory, such considerations are of critical importance to the successful conclusion of this project.

3.4 Scratch Data and I/O:

Scratch storage of 10 TB is presently sufficient to perform simulations and data analysis. Since restart requires an I/O of almost the entire memory (~ 20 TB), it does not appear to be cost efficient at present. While an increase by a factor of 10–20 in scratch storage will be needed to analyze larger lattices, data limitations due to I/O bandwidth are not expected to cause worrisome performance limitations.

3.5 Long-term and Shared Online Data:

Simulation output needs to be stored for a time scale of about 2–3 years, mainly for comparison with future simulations and re-analysis. Overall storage needs are comparable to scratch requirements, and will be on the order of 10s of TB. Since we anticipate computations to be performed at different computing facilities, it is highly desirable to have a central place for data collection, which can be accessed by all collaborators for analysis purposes.

3.6 Archival Data Storage:

It is not clear that archival storage beyond a 2–3 year time scale will be needed. Over this time frame, storage of 10s of TBs will be required.

3.7 Workflows:

Simulations are presently performed at the NERSC computing facility and create the output of physical observables. Data analyses are performed subsequently, involving simple C++ programs or shell scripts with short run-times, and are carried out either directly at NERSC or by downloading data to local machines. Visualization is mostly performed on local machines, after downloading relevant data. With increasing amounts of data, in the future we expect to carry out a major part of analysis and

visualization directly at the computing facilities. Sufficiently fast **ssh** access to the login nodes will be required to perform these tasks efficiently.

3.8 Many-Core and/or GPU Readiness:

At present, this software suite runs on conventional CPUs only and does not take advantage of GPUs or Many-Cores architectures. Currently, funding is not available in this area to support a workforce in support of porting codes to heterogeneous architectures.

3.9 Software Applications, Libraries, and Tools:

Standard C++ compilers and libraries such as GSL, FFTW and LAPACK are sufficient for our simulations.

3.10 HPC Services:

Advanced training in software optimization for future architectures and 3D data visualization will be required.

3.11 Additional Needs:

A sufficient level of effort is of critical importance for this project as software maintenance, executing jobs, and performing data analysis and visualization will require a significant amount of time. In view of the increasing complexity of these calculations, it would be highly desirable to have a full-time person dedicated to the computational aspects of this project.

Requirements Summary Worksheet

Non-Equilibrium QCD Code: CYM+WilsonSolver	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	5M	1,500M	30B
Computational node hours (homogeneous many-core)	N/A	N/A	N/A
Memory per node	16 GB	8–16 GB	4–16 GB
Aggregate memory	20 TB	2,000 TB	5,000 TB
Data read and written per run	0.05 TB	0.25 TB	0.5 TB
Maximum I/O bandwidth needed	400 GB/sec	800 GB/sec	800 GB/sec
Percentage of runtime for I/O	<1%	<1%	<1%
Scratch file system space needed	10 TB	100 TB	200 TB
Permanent online data storage	3 TB	10 TB	20 TB
Archival data storage needed	5 TB	5 TB	5 TB

Phases and Properties of Hot-dense Strongly Interacting Matter

Lead Authors: Swagato Mukherjee and Patrick Steinbrecher (Brookhaven National Laboratory)

1. Description of Research

1.1 Overview and Context:

A central goal of the current physics program at the Relativistic Heavy Ion Collider (RHIC) of the Brookhaven National Laboratory (BNL) is the exploration of the phase diagram of strongly interacting matter as articulated in the 2015 NSAC Long Range Plan. In particular, to map out the phase diagram and to determine the properties of baryon-rich QCD matter at high temperatures, a major experimental effort, the RHIC Beam Energy Scan (BES) program, is now underway. A series of measurements of heavy-ion collisions spanning the collision energy can explore the properties of a large region of the QCD phase diagram. In addition to these high precision experimental measurements, a comprehensive understanding of the phase and properties of hot-dense QCD matter requires a framework for modeling the salient features of heavy-ion collisions at BES energies so that the experimental results can be quantitatively described. The crucial aspect of this program is embedding the various bulk and transport properties of the underlying fundamental theory, i.e. equilibrium QCD inputs, into a dynamical scheme. At present lattice QCD is the only viable technique for approximation-free, parameter-free *ab-initio* calculations of the equilibrium properties of hot-dense QCD matter.

1.2 Research Objectives for the Next Decade:

The goal is nonperturbative, parameter-free determinations of the bulk thermodynamic properties, as well as the phase structure of hot-dense strongly interacting matter from the fundamental theory, using state-of-the-art lattice QCD supercomputing. The important physics objectives for the next decade are the following:

- Computations of the QCD equation-of-state at non-zero baryon densities.
- Calculations of non-Gaussian fluctuations of various conserved charges.
- Determination of the freeze-out parameters for heavy-ion collisions.
- Exploration of the thermodynamic signatures that experimentally uncharted hadrons might produce in RHIC's BES.
- Determination of the QCD phase boundary in the temperature and baryon chemical potential plane and constraining the location of the QCD critical point.

2. Computational and Data Strategies

2.1 Approach:

Lattice QCD is an alternative regularization of QCD suitable for large-scale numerical computations. In this formulation the four-dimensional space time is discretized, with a finite numbers of points along the space and the (Euclidean) time directions. Quarks are defined on the lattice sites and the gluons on the connecting links between the sites. Interactions between quarks and gluons are implemented in accordance with the discretized version of the QCD Lagrangian. Temperature is introduced by breaking the Lorentz symmetry by allocating different numbers of points along the space and the time directions.

An imbalance of baryonic charge can be introduced by coupling the corresponding chemical potential to the net baryon charge density added to the QCD Lagrangian. The QCD path integral is represented as an ordinary integral over all the site and link variables. The above-mentioned large dimensional integral is numerically evaluated by stochastic Monte Carlo sampling. The standard Monte Carlo sampling employed during lattice QCD calculations is inapplicable for a non-zero value of the baryon chemical potential, as in this case the sampling weight, i.e. the QCD action, becomes a complex number and thus cannot be interpreted as a probability density. At present, the most successful method for undertaking lattice QCD calculations at non-zero baryon chemical potentials for QCD with physical quark masses and close to the continuum and thermodynamic limit is by using the well-established Taylor expansion method. Herein, one expands the QCD partition function, or equivalently the pressure, in a power series of the chemical potential around the vanishing values of the chemical potential. The coefficients of this series are known as the generalized susceptibilities associated with the specific conserved charge. Since these generalized susceptibilities are defined at vanishing chemical potentials, the usual lattice QCD simulations can be used to compute these coefficients.

2.2 Codes and Algorithms:

Lattice QCD computations of up to $6 \cdot 10^{\text{th}}$ order of conserved charge fluctuations is a two-step process: (i) generation of about 100K independent gauge field configurations; (ii) measurement of 10 separate operators involving the Dirac fermion matrix for each of those background gauge-field configurations. The gauge-field configurations will be generated by using the Rational Hybrid Monte Carlo. For each gauge configuration one requires $O(10^6)$ inversions of the Dirac matrix, a very large sparse matrix, to compute the relevant fermionic operators for the conserved charge fluctuations. Both gauge configuration generation and measurements are dominated by the inversions of the Dirac matrix, and these inversions are usually carried out using the Conjugate Gradient algorithm.

3. Current and Future HPC Needs

3.1 Computational Hours:

Presently the calculations are performed on GPUs and require about $\sim 7 \times 10^6$ hours on K20x cards or equivalently $\sim 10^9$ conventional core hours (using USQCD conversion). Present calculations are performed on lattices with temporal extent $N_t=8$. To control discretization effects and eventually perform continuum extrapolations calculations with $N_t=12$ and $N_t=16$ will be needed. The cost of calculations increases as N_t^4 (simple volume scaling). Furthermore, higher statistics are needed when performing calculations at larger N_t . We estimate that the statistics have to be increased by at least factor of 4 and 20 for $N_t=12$ and $N_t=16$, respectively. Therefore, the computing needs for this area will be approximately 20x larger in 2020 and approximately 320x larger in 2025.

3.2 Parallelism:

We have implementations of our algorithms in CUDA for the GPU and in a combination of OpenMP and intrinsics for the CPU, i.e. we take full advantage of the on-chip parallelism. Both frameworks are designed in an architecture independent approach, thus, can be extended to future hardware generations. Our main computation naturally has a high degree of parallelism. We are able to assign many inversions of a large, sparse matrix independently to each node. Communication in our program is limited to the start-up phase which is insignificant for the total run-time. By doing so, we can scale up to 20000 nodes without any loss in efficiency, i.e. we maintain our fast single node performance. If in the future our problem does not fit into on-chip memory, we plan to assign a single matrix inversion to a few nodes but still keep our total job size large in order to reduce the time to solution.

3.3 Memory:

Presently, the memory requirement per node is approximately 16 GB. In the future, we expect this to increase up to approximately 32 GB in 2020 and approximately 256 GB in 2025. To keep our time to solution at a reasonable level, we need to run jobs using at least with approximately 5,000 nodes. This sets the minimum requirement of the aggregate memory.

3.4 Scratch Data and I/O:

Our present need is approximately 20 TB, and the expected future need is 40 TB in 2020 and 200 TB in 2025.

3.5 Long-term and Shared Online Data:

Our present need is approximately 300 TB and the expected future need is 2,500 TB in 2020 and 24000 TB in 2025.

3.6 Archival Data Storage:

Presently, we require approximately 100 TB, and the expected future need is 700 TB in 2020 and ~ 4,000 TB in 2025.

3.7 Workflows:

Our present sequential workflow is (i) gauge field configurations, (ii) calculations of eigenvectors of the Dirac matrix on these configurations, (iii) measurements of observables on these configurations, with the help of the pre-computed eigenvectors.

3.8 Many-Core and/or GPU Readiness:

Our algorithms naturally have very good locality (i.e., they can stay for some time in even small caches), and therefore deepening memory hierarchies will certainly be helpful in overcoming the memory-bandwidth gap. In particular, our cache-oblivious algorithms use space-filling curves to improve spatial locality. For deepening cache hierarchies, we need to further optimize these space-filling curves. For the latter, it would be helpful to have a direct programmable cache rather than a cache managed by some protocol.

3.9 Software Applications, Libraries, and Tools:

The codes are written using C++, CUDA, OpenMP, and MPI using standard compilers and libraries.

3.10 HPC Services:

Transferring large sets of gauge configurations between different computing facilities might become an issue in the future and may demand fast file transfer connectivity between the facilities.

3.11 Additional Needs:

At present, we do not anticipate any additional needs.

Requirements Summary Worksheet

Phases of Hot and Dense Matter Code: Finite T and μ_B Lattice QCD	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	1B	20B	600B
Computational node hours (homogeneous many-core)	7M	140M	320M
Computational node hours (with GPU or accelerator)	7M	140M	4,200M
Memory per node	16 GB	32 GB	256 GB
Aggregate memory	70 TB	130 TB	1,000 TB
Data read and written per run	20 TB	40 TB	200 TB
Maximum I/O bandwidth needed	400 GB/sec	800 GB/sec	4,000 GB/sec
Percentage of runtime for I/O	<0.005%	<1%	<1%
Scratch file system space needed	500 TB	4,000 TB	40,000 TB
Permanent online data storage	300 TB	2,500 TB	24,000 TB
Archival data storage needed	100 TB	700 TB	4,000 TB

Spectral Functions, Transport Coefficients and Hadron Properties from Lattice QCD

Lead Author: Olaf Kaczmarek (Bielefeld University)

1. Description of Research

1.1 Overview and Context:

The in-medium modifications of hadrons as well as transport coefficients are important properties of strongly interacting matter that are encoded in spectral functions of the corresponding hadronic observables. These properties are described by Quantum Chromodynamics (QCD), the theory of strong interaction physics. Lattice QCD calculations provide the only *ab-initio* method known today to extract these properties from calculations of hadronic correlation functions. High-end computing is mandatory for such calculations. In the light quark sector these correlation functions reflect the temperature dependence of chiral symmetry and confinement properties of matter and provide access to experimentally relevant probes, thermal dileptons, and thermal photons. The dissociation of heavy quark bound states in the quark gluon plasma and the melting temperatures of different states are important properties for the understanding of strongly interacting matter at high temperatures and densities as it is created in heavy-ion experiments at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory and the Large Hadron Collider (LHC) at CERN in Switzerland. In addition, information on transport coefficients, like the electrical conductivity and heavy quark diffusion coefficients, can be extracted from the low-frequency behavior of the spectral functions. These provide crucial information about the matter produced in heavy-ion experiments and are needed for the modeling of its hydrodynamic expansion.

1.2 Research Objectives for the Next Decade:

The determination of transport and spectral properties of strongly interacting matter, until now, has been mainly limited to the quenched approximation in which effects resulting from dynamical quark–antiquark pair creation are suppressed. In the next decade, it will become important to extend these calculations to cases where dynamical quark degrees of freedom are included. These full QCD calculations are especially important in the confined phase and in the transition region from hadronic matter to the quark-gluon plasma, where the effect of light quark degrees becomes more prominent, e.g. leading to a change of the nature of the transition. Also, in the heavy-quark sector light-quark degrees of freedom will play an important role, not only for heavy-quark bound states but also for open-flavor hadrons, and for a quantitative determination of the dissociation temperatures and transport coefficients.

2. Computational and Data Strategies

2.1 Approach:

Lattice QCD calculations in this project require large and fine lattices, which currently limit calculations to the quenched approximation. In the next decade it will become important to extend these calculations to the case of full QCD, including dynamical light-quark degrees of freedom. While the calculation of the hadronic correlation functions is rather straightforward and possible with currently available high-performance computing resources, the generation of gauge-field configurations in full

QCD including physical light-quark degrees of freedom requires computing resources that are orders of magnitude larger than those currently used in quenched QCD studies.

2.2 Codes and Algorithms:

We have developed lattice QCD codes over the past decades that are constantly ported and optimized for various HPC architectures. The work flow is usually split into two parts, one for the generation of the gauge-field configurations and another one for the measurement of observables. For most applications, the generation of the gauge-field configurations is the most time-consuming part, especially when approaching physical quark masses. Here various algorithmic improvements are used in the Hybrid Monte Carlo evolution. Together with increasing computational resources, it is now possible to use physical quarks masses for finite temperature lattices up to $64^3 \times 16$ and finer lattices, required for continuum extrapolations, in the next decade. For the measurement part, e.g., the calculation of hadronic correlation functions where the large fermion matrix needs to be inverted, inversion methods are used, like conjugate gradient or multi-mass conjugate gradient solvers. Also algorithmic improvements can be used, e.g. multigrid, deflation or block solvers.

3. Current and Future HPC Needs

3.1 Computational Hours:

Within the quenched approximation, the generation of the gauge-field configurations and the calculation of the hadronic correlation functions require approximately 100M to 200M core hours on conventional cores. The computational requirements for full QCD calculations will increase by 1 to 2 orders of magnitude in the next decade.

3.2 Parallelism:

The codes in use are parallelized using MPI and can use OpenMP for node-level parallelization. Optionally GPUs can also be used as accelerators.

3.3 Memory:

The memory requirements vary depending on the lattice sizes from a few GB to a few TB. These requirements will not change in the coming decade.

3.4 Scratch Data and I/O:

Scratch data is required for the gauge-field configurations on the order of a few TB. The main I/O requirements arise from reading and storing the gauge-field configurations at the beginning and end of each job. Therefore, the percentage of the total runtime devoted to I/O is usually small for reasonable I/O bandwidths.

3.5 Long-term and Shared Online Data:

The generated gauge-field configurations can be used for various physics topics, and a long-term storage is important to allow for future calculations. Sharing these configurations for other projects and other collaborations maximizes the scientific impact.

3.6 Archival Data Storage:

The required amount of archived gauge-field configurations is on the order of a few PB.

3.7 Workflows:

In the first step, gauge-field configurations are generated, and in a later stage, observables like hadronic correlation functions are calculated in different runs.

3.8 Many-Core and/or GPU Readiness:

The programs used in this study can utilize many-core as well as GPU resources. The codes are optimized for various architectures including different accelerators.

3.9 Software Applications, Libraries, and Tools:

No special software applications, libraries, or tools are needed. The code is written using standard compilers and libraries.

3.10 HPC Services:

No special services are required.

3.11 Additional Needs:

We do not anticipate any additional needs at this time.

Requirements Summary Worksheet

Spectral Functions and Transport Coefficients Code: Lattice QCD spectral functions	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	100M-+200M	800M	2,000M
Computational node hours (homogeneous many-core)	1M-3M	12M	30M
Memory per node	64 GB	64 GB	64 GB
Aggregate memory	20 TB	20 TB	20 TB
Data read and written per run	2 TB	2 TB	2 TB
Maximum I/O bandwidth needed	400 GB/sec	800 GB/sec	4000 GB/sec
Percentage of runtime for I/O	2%	2%	2%
Scratch file system space needed	20 TB	20 TB	20 TB
Permanent online data storage	20 TB	20 TB	20 TB
Archival data storage needed	1 PB	2 PB	3 PB

Modeling Heavy-Ion Collisions

Lead Authors: Scott Pratt (Michigan State University) and Björn Schenke (Brookhaven National Laboratory)

1. Description of Research

1.1 Overview and Context:

Heavy-ion collisions can create novel states of matter, but inferring bulk properties from these ephemeral conditions requires careful comparison of complex multi-component models of the collision dynamics to large-scale heterogeneous data sets from both the Large Hadron Collider (LHC) and the Relativistic Heavy-Ion Collider (RHIC). A few dozen unknown parameters are required to encapsulate our uncertainties in the collision and the bulk properties of the quark-gluon plasma. Rigorously constraining these parameters allows quantitative conclusions regarding the equation of state, viscosity, chemistry, and phenomenology of the environments created in heavy-ion collisions to be drawn.

Sophisticated three-dimensional models of the evolution are numerically intensive and require significant computational resources once one explores the high-dimension parameter space and considers fluctuating geometries. This is necessary if the field is to achieve its goal of quantitatively obtaining the properties of the quark-gluon plasma from experiments. In addition to determining properties, this effort would be pivotal in addressing the question of whether or not a critical point exists in the QCD phase diagram. U.S. funding agencies will have invested over three billion dollars in the RHIC and LHC heavy ion programs by 2020, and without this final effort to rigorously compare models and data, principal goals of the programs will go unrealized.

1.2 Research Objectives for the Next Decade:

The upcoming beam energy scan at RHIC provides great challenges. By lowering the beam energy matter is explored at increasingly high baryon density. Compared to analyses of LHC and highest-energy RHIC data, computational resource requirements increase dramatically. An order of magnitude more data is involved by considering results from multiple beam energies and from multiple target projectile combinations. More importantly, high-energy calculations invoked boost-invariant symmetries justifying two-dimensional hydrodynamic treatments. Fully three dimensional hydrodynamic calculations, which increase the computational resource requirements by at least two orders of magnitude, are unavoidable at lower energy.

2. Computational and Data Strategies

2.1 Approach:

Models incorporate parametric descriptions of the initial state or fully dynamic pre-equilibrium models, relativistic viscous hydrodynamics, microscopic simulations of hadronic scattering for the breakup stage, plus a number of modules for calculating correlations of various sorts or for interfacing the aforementioned model elements. Running a model for a single point in the multi-dimensional parameter space involves addressing numerous beam energies, several projectile target combinations, and the entire range of impact parameters for each case. This might require approximately 10^5 CPU days to create enough data and compare it to the data sets that should exist by 2020. The output of one such calculation is of the order of one TB for one point in parameter space. Using modern emulator methods,

$>10^3$ different points are required to explore the high-dimension parameter space. This requires approximately 1 Billion CPU core hours and results in petascale output.

2.2 Codes and Algorithms:

Current Workflow: 1) Run $N \sim 10^3$ full model runs using semi-random parameters. 2) Determine principal components of final-state observables. 3) Build and design an emulator, tool that interpolates principal components from the initial N full-model runs. 4) Perform Markov Chain comparing principal components to experimental measurements 5) Augment N runs with a few dozen more full-model runs in like region of parameter space 6) Repeat 3-6 until convergence.

Only steps 1 and 5 involve large-scale computing. Thus far, this method has only been applied for highest-energy data where the hydrodynamic codes are two-dimensional and numerical costs are greatly reduced.

3. Current and Future HPC Needs

3.1 Computational Hours:

Three-dimensional hydrodynamic codes represent the great bulk of future computing needs. Such codes typically use ≤ 2 GB of memory per core and run for the better part of one day on a commodity CPU. The calculation must be repeated 10^8 times to account for fluctuating initial conditions, varying beam energies, and projectile combinations, plus sampling parameter space.

3.2 Parallelism:

Current hydrodynamic codes use multi-node parallelism. Typically, we use around 100 cores per event. Between different events, communication is not needed. Usually the grid is split over cores only in one dimension. It is conceivable that future codes will parallelize in three dimensions, leading up to one million cores per event. These codes are good candidates for GPU architectures given the locality inherent to hydrodynamic treatments.

3.3 Memory:

The hydrodynamic codes employ meshes of $\sim 200 \times 200 \times 200$ and typically need approximately 5 GB of main memory.

3.4 Scratch Data and I/O:

Current data storage is on the order of 10s of TBs. This can increase significantly when storing initial state data and freeze out surfaces as well as hydro evolution data that can be used in other simulations. We estimate the scale of the output to be on the order of 10s to 100s of PB.

3.5 Long-term and Shared Online Data:

For recent analyses, output needs to be readily available for 3–4 years. Recent analyses have only a few TB of output, but the beam energy scan, combined with the increased data requirements for fluctuating initial conditions, will increase this need by more than two orders of magnitude.

3.6 Archival Data Storage:

It is not clear whether it is cost-efficient to store large amounts of data beyond 3–4 years. The hydrodynamic hyper-surfaces should be stored long term, but that is less than one percent of the output. Hadronic simulations produce 99% of the output, but require only ~1% of the CPU. If needed, the final output could be regenerated from the hyper-surfaces at low cost. When the complete hydrodynamic evolution is to be saved, the storage requirement will increase significantly.

3.7 Workflows:

For more complex models initial state calculations need to be run and their output stored. Then the hydrodynamic evolution in 3+1D is performed with the freeze-out surfaces as output. Using these files, the thermal spectra are computed and interfaced to the hadronic transport model. Finally, many Monte Carlo simulations per hydro event of the hadronic transport including resonance decays are performed.

3.8 Many-Core and/or GPU Readiness:

Codes are not yet GPU-ready, but hydrodynamic treatments may be good candidates for such efforts. Some GPU-based hydro codes are already being developed.

3.9 Software Applications, Libraries, and Tools:

All work can be completed with standard C++ compilers and libraries.

3.10 HPC Services:

Advanced training for porting hydrodynamic codes onto GPUs could be valuable, as could training into maintaining large data sets.

3.11 Additional Needs:

As mentioned above, this work requires a sufficient level of effort. One full-time person, whose main function would be to facilitate the computation and collaboration of the effort, would be essential.

This case study describes a single effort at global analysis. Because of needing to run for approximately 10^3 points in parameter space, it is likely to require the bulk of resources needed in this area, but an additional factor of 2 of resources could be required across the field. Sophisticated microscopic models of the initial state also require resources even if they are not explicitly incorporated into the global analysis. A half dozen groups are involved in developing hydrodynamic codes, parts of which might be used in this case study. Each group may need resources on the scale of 5% of the resources described here.

Requirements Summary Worksheet

Code: Modeling HIC	Column 1: Current Usage	Future Usage: 2020 (As a factor of column 1)	Future Usage: 2025 (As a factor of column 1)
Computational core hours (conventional)	10M	1B	1B
Computational node hours (homogeneous many-core)	N/A	N/A	N/A
Memory per node	5 GB	5 GB	5 GB
Aggregate memory	5 GB	5 GB	5 GB
Data read and written per run	10 TB	1,000 TB	1,000 TB
Maximum I/O bandwidth needed	10 GB/sec	10 GB/sec	10 GB/sec
Percentage of runtime for I/O	5%	5%	5%
Scratch file system space needed	100 TB	,1000 TB	10,000 TB
Permanent online data storage	10 TB	100 TB	1,000 TB
Archival data storage needed	10 TB	100 TB	100 TB

