Computational Nuclear Physics Meeting

SURA Headquarters, Washington DC, July 14-15, 2014

REPORT

Prepared by the Computational Nuclear Physics Meeting Writing Committee

A. Burrows, J. Carlson, W. Detmold, R. Edwards, R. Furnstahl, W. Haxton, W. Hix, F. Karsch, W. Nazarewicz, P. Petreczky, D. Richards, M.J. Savage.

Contents

Chapter 1

Executive Summary

Large scale numerical calculations in theoretical nuclear physics play an increasingly prominent role in our understanding of the physics of strongly interacting systems from hadrons and nuclei to hot and dense matter. These calculations not only impact our theoretical understanding but also guide present and future experimental programs. The importance of High Performance Computing in nuclear physics was recognized in the 2007 Long Range plan and reinforced in the vision presented in the recent National Academy of Sciences (NAS) report "Exploring the Heart of Matter":

> *High performance computing provides answers to questions that neither experiment nor analytic theory can address; hence, it becomes a third leg supporting the field of nuclear physics.*

and its consequent recommendation:

> *A plan should be developed within the theoretical community and enabled by the appropriate sponsors that permits forefront computing resources to be exploited by nuclear science researchers, and establishes the infrastructure and collaborations needed to take advantage of exascale capabilities as they become available.*

An article in HPCwire on *Meet the Exascale Apps* identifies potential application areas for exascale development. It notes the emergence of nuclear physics as a driver that was not evident in the 1990s (Fig. 2 therein).

Computational Nuclear Physics Meeting

In April 2014, the DOE and NSF charged NSAC to survey the opportunities and priorities for US nuclear physics research, and to recommend a long range plan (LRP) to advance that research over the next ten years. In light of the integral role of computational nuclear physics across the nuclear physics program, a Computational Nuclear Physics Meeting was held at the headquarters of the Southeastern Universities Research Association in Washington DC on July 14-15, 2014. The aim of the meeting was to review the status, objectives and opportunities for computational nuclear physics, and to determine the human and computational resources for computational nuclear physics needed to advance nuclear-physics research over the next ten years. The outcome of the meeting is presented here.

The meeting attracted around 40 attendees encompassing domain scientists representing the major thrusts of computational nuclear physics, and representatives of the DOE and NSF. The first day comprised a review of progress and activities since the last LRP in 2007, a discussion of the input that would be needed for the LRP, and finally the development of a strategy for computational nuclear physics. The second day was devoted to a discussion of the human and computational resources needed, before concluding at noon with a unanimous recommendation and resulting request.

RECOMMENDATION AND REQUEST

- **Recommendation:** Realizing the scientific potential of current and future experiments demands large-scale computations in nuclear theory that exploit the US leadership in high-performance computing. Capitalizing on the pre-exascale systems of 2017 and beyond requires significant new investments in people, advanced software, and complementary capacity computing directed toward nuclear theory.

- **Request:** To this end, we ask the Long-Range Plan to endorse the creation of an NSAC subcommittee to to develop a strategic plan for a diverse program of new investments in computational nuclear theory. We expect this program to include:

  - new investments in SciDAC and complementary efforts needed to maximize the impact of the experimental program;
  - development of a multi-disciplinary workforce in computational nuclear theory;
  - deployment of the necessary capacity computing to fully exploit the nations leadership-class computers;

  with support ramping up over five years towards a level of around $10M per annum.

Science

Below we describe some of the activities that contribute to the vision for the field presented at the meeting, and the computational requisites to satisfy it. The details pertaining to the major computational efforts are contained in Chapter 2.

Hadrons, Nuclei and Nuclear Matter

Quantum Chromodynamics (QCD) is the underlying theory of the strong interactions that, together with the electroweak force, describes the structure of nuclei, the nature of their constituent hadrons, and the phases of strongly interacting matter. High-impact nuclear-physics experiments will elucidate QCD, test fundamental symmetries and probe novel interactions in and beyond the standard model at an unprecedented level. Equally precise theoretical analyses of QCD are necessary to interpret these experiments. The powerful numerical technique of lattice gauge theory enables key properties of the world around us to be computed from first principles. Thus, lattice QCD calculations can determine the bound states of the theory, and describe how the fundamental quarks and gluons of QCD give rise to the observed protons, neutrons, pions, and the other hadrons. They can determine how charge, current, and matter are distributed within a hadron, and contribute to the building of a three-dimensional picture of the proton. The emergence of the nuclear force from QCD can be investigated, leading to a refinement of the chiral nuclear forces, and first-principle calculations of the structure and reactions of light nuclei can be performed. Precision calculations of QCD will enable careful probes of the fundamental symmetries of nature. The phase structure of strongly interacting matter can be investigated, describing how the relevant degrees of freedom of QCD change under conditions of high temperature and/or high density, and the role these phases play in the cosmos can be revealed.

Lattice calculations have a key role both in complementing and supporting future experiments, and in fully capitalizing on the current and future DOE experimental nuclear physics programs. Lattice QCD calculations will predict the spectrum and properties of so-called exotic mesons, in which the gluonic degrees of freedom are manifestly exposed, that are the target of the GlueX experiment at the 12 GeV upgrade of Jefferson Laboratory. Calculations of nucleon form factors, generalized parton distributions, and transverse-momentum-dependent distributions will provide a more complete three-dimensional tomography of the nucleon than can the experimental programs at JLab and at RHIC-spin provide alone. Nucleon and nuclear matrix elements are essential for interpreting precision measurements of hadronic parity violation and neutron and nuclear beta decay, as well as experimental limits established on electric dipole moments (EDMs).

Understanding the properties of strongly interacting matter created in heavy ion collisions at RHIC and LHC also requires large scale computations. Heavy ion collisions create a very complex dynamically evolving systems. At very early times the system created in the collisions can be described by classical-statistical simulations of dense gluon fields. At later stages when the system is thermalized its properties, such as the transition temperature, equation of state, fluctuations of conserved charges, spectral functions etc can be studied by using first principle lattice QCD calculations. Finally, for exploiting the wealth of experimental data obtained in heavy ion collisions a comprehensive model-to-data comparison is needed involving 3+1 hydrodynamic calculations for variety of initial conditions and for many values of input physical parameters.

Finally, lattice studies of the nuclear and hypernuclear forces will refine existing nuclear forces, particularly the multi-nucleon interactions, and provide a first-principles QCD underpinning to studies of nuclear structure and reactions and to nuclear astrophysics, both areas which are central to the FRIB experimental program, and provide data for key interactions, such as those of hypernuclei, for which there is a paucity of empirical data.

Nuclear Structure and Reactions

Large-scale nuclear physics computations are dramatically increasing our understanding of nuclear structure and reactions, and of the properties of nuclear matter. In light nuclei, *ab initio* calculations of reactions, including those for fission and electroweak processes, are being used to systematically reduce uncertainties and to predict experimentally inaccessible data and processes. *Ab initio* approaches are being extended to reactions on medium-mass nuclei. Density Functional Theory and its extensions are being applied to medium-mass and heavy nuclei, including neutron-rich systems, nuclear matter, and the structure of neutron star crust. The suite of computational methods, including Quantum Monte Carlo, Configuration Interaction, Coupled Cluster, and Density Functional Theory, now scale efficiently to the largest computers.

These studies will have a direct impact on the experimental nuclear physics program. Computational studies of the strongly correlated matter found in nuclei and neutron stars impact the current experimental programs at NSCL and ATLAS, and guide the future program at FRIB. Computations of the neutron distribution in nuclei and of electroweak process are key to the interpretation of results at JLab. Simulations of light-ion thermonuclear reactions and fission are relevant to NNSA priorities. Precise calculations of nuclear matrix elements provide crucial input for the interpretation of fundamental interaction experiments. Finally, accurate solutions of the strongly interacting quantum many-body systems will yield new insights and the ability to calculate phenomena, processes, and states of matter that are difficult or impossible to measure experimentally, such as the crust of a neutron star or the core of a fission reactor.

Nuclear Astrophysics

A central goal of nuclear physics is the detailed explanation of the origin of the elements and of their isotopes. Overwhelmingly, these elements are produced in stars, either during quiescent thermonuclear burning stages, or explosively, as in core-collapse and thermonuclear supernovae. Most of the elements up to the "iron peak", named for the prominent peak in the solar abundances centered on iron, are ejected in supernova explosions, and an understanding of stellar explosions and stellar evolution is central to an understanding of the abundance pattern of the nuclei around us. Stellar evolution calculations involve both nuclear ratio rates generated either theoretically or through experiment, and three-dimensional turbulence and magnetic interactions. Stellar explosions are always multi-dimensional requiring state-of-the-art radiation/hydrodynamical simulations. As a consequence, addressing this key nuclear physics goal of the origin of the elements entails the sophisticated numerical simulations employing the latest computational tools and the exploiting the most powerful supercomputers.

Computational nuclear physics is essential to capitalize on the DOE Office of Nuclear Physics experimental programs. The astrophysics of neutron-rich nuclei is one of four scientific pillars of FRIB. Comparisons of supernova calculations with flagship experimental observations at the Intensity Frontier of the DOE High Energy Physics (HEP) program could reveal new physics beyond the Standard Model. Finally, investigations of the nuclear equation of state can be compared with experiments such as those at JLab aimed at measuring the weak charge radius of lead and calcium.

Neutrinos and Fundamental Symmetries

Neutrinos and fundamental symmetries are an increasingly important part of the national nuclear physics program, addressing very basic questions regarding neutrino properties (mixing angles, the absolute mass scale and the mass hierarchy), the nature of dark matter, and physics beyond the Standard Model. Computational nuclear physics plays an important role in fundamental symmetries and neutrino physics at both the hadronic and the nuclear level, often being essential in turning measured rates into constraints on a fundamental interaction. The relevant calculations span much of the effort discussed above and in the coming years, advances in computational nuclear physics will be central to optimizing the impact of many of the experiments in this domain.

Neutrinos and weak interactions are critical pieces of many astrophysical environments, including core-collapse supernovae. In addition to scattering from nuclei, in the very high flux supernovae environment coherent neutrino-neutrino scattering may be important in nucleosynthesis and observed neutrino fluxes. Terrestrial neutrino experiments at reactors and accelerators depend in a crucial way on our understanding of neutrino interactions with nuclei, and well-controlled nuclear theory is needed to maximize the impact of these experiments. In double-beta decay, improved calculations of the nuclear matrix element will impact both the size of experiment necessary to cover the inverted mass hierarchy region and the extraction of the absolute mass scale from a measured rate.

In many cases, fundamental symmetry experiments have standard model contributions that must be understood through a combination of lattice QCD and nuclear theory. Neutron beta decay experiments, searches for neutron, nuclear and atomic EDMs, future $\mu \to e$ conversion experiments, and terrestrial searches for dark matter are all predicated on understanding these contributions. Advances in computational nuclear theory will enable the necessary hadronic and nuclear contributions to be determined with quantified uncertainties on relevant timescales.

Computational Resources

This ambitious program in computational nuclear physics requires substantial computational resources. Advanced calculations that can resolve the key issues of nuclear physics require leadership computing at the largest scales (capability computing), medium scales (capacity computing), and a software infrastructure to effectively exploit that computing. The nuclear physics community has aggressively pursued these resources and thereby received substantial allocations on leadership-class facilities such as those under DOE's INCITE program, and programs within the NSF and NNSA. It is through the generous access to such large scale computational resources, and the recognition of the need for a range of computational capabilities to exploit them, such as the dedicated facilities of USQCD, that such impressive progress has been made. However, as outlined in the 2009 Nuclear Physics Exascale Report, substantially more computational resources will be needed to capitalize on the scientific potential made possible by computational nuclear physics.

Over the last year, on the order of 1 Billion CPU core-hours have been made available to all of theoretical computational nuclear physics within the U.S., including QCD, nuclear structure, and nuclear astrophysics. This corresponds to about 0.1 Petaflop-years which places this field at the earliest part of the time-line described in the 2009 Nuclear Physics Exascale Report. New generations of calculations will require substantially more computing time as the fidelity of the simulations improves. As discussed below, major investments in computer science and applied mathematics manpower are needed to maintain this momentum. These investments will be crucial in transitions to new computational architectures and programming models beyond the petascale.

Workforce, Education, Funding and Career Paths

Computational nuclear physics bridges many areas of science, and as such, the expertise of a broad range of individuals including physicists, computer scientists, applied mathematicians is vital to the success of the program. Contributions at all levels from students to senior researchers are also important to the large scale efforts that are necessary. The 2013 Tribble Committee Report on Implementing the 2007 Long Range Plan emphasized the role of junior scientists:

> "People remain the key factor. In particular, early-career scientists working at the interface between nuclear theory, computer science, and applied mathematics are critical to make future impact, especially in the era of extreme computing that demands the novel coding paradigms and algorithmic developments required by novel architectures."

Even with the raw computational power of national computing growing almost exponentially, the overall improvement in NP computing since the last long range plan far outstrips this growth alone. The ingenuity and dedication of the NP workforce, and collaboration with colleagues in computer science and applied mathematics, have led to fundamentally better algorithms and structurally more efficient code bases, allowing computations that at the time of the last long range plan we would not have thought possible over such a short timescale. As computational environments become more diverse and more specialized, it is crucial that the workforce include those individuals with skills sufficient to master these challenges which extend beyond the traditional purview of nuclear theory. This workforce must follow the emerging trends in these architectures and will develop the algorithms and software infrastructure to enable their exploitation and advance the overall program. To maintain this progress in light of the demands on large-scale computation from the NP program, new investment will be essential to develop the needed multi-disciplinary workforce in computational nuclear theory.

The interdisciplinary SciDAC program has been essential across several areas of nuclear physics, but important activities remain unsupported. *The strengthening and broadening of SciDAC and complementary programs will be essential if scientific goals laid out here are to be realized.*

Many of the innovations and a significant amount of the burden of performing the calculations has been borne by junior researchers: graduate students, postdoctoral scholars, and tenure-stream faculty. The quality of these junior scientists throughout the nuclear physics subfields is outstanding. Over the last 7 year period since the last previous LRP exercise, many junior faculty and permanent lab positions have been filled by computationally oriented nuclear theorists, both within the US and abroad. These junior researchers are very successful in acquiring competitive funding for their research, winning numerous early career awards and prestigious prizes. Notably, of the fourteen DOE Early Career research awards in nuclear theory since the programs inception in 2010, ten have been awarded to researchers with a strong focus on computational nuclear theory. These successes show the field is vibrant and ready to address the challenges discussed above.

More broadly, nuclear physics plays an important role in training computational scientists to help meet the needs of the workforces of the DOE and of industry. Further strengthening the computational effort in NP will contribute to the broader DOE mission and to society as a whole, where large scale computation and data-analysis are becoming increasingly important skills.

Chapter 2

Science Opportunities

In this chapter, we discuss the main physics missions addressed through computational nuclear theory in more detail and highlight computational and workforce aspects required for the science goals. Fundamental symmetries and neutrinos impact, and are impacted by, all three main thrusts, lattice QCD, nuclear structure and nuclear astrophysics and our discussions are broken up accordingly.

2.1 Hadrons, Nuclei and Nuclear Matter

The properties and interactions of strongly interacting matter in a stationary state can be determined from first principles with numerical calculations of discretized versions of Quantum Chromodynamics, called lattice QCD. Deviations in observables from their QCD values, introduced by the discretization and the finite spacetime volume, can be systematically reduced and quantified using well-established field theory techniques. Lattice QCD allows for the calculations of the hadronic spectrum, nucleon properties and few nucleon interactions, as well as the structure and properties of light nuclei. These calculations, referred to as "Cold QCD", are relevant for the experimental programs at existing and future facilities such as JLab, FAIR, J-PARC and FRIB and also in experiments seeking to probe fundamental symmetries and look for physics beyond the Standard Model in hadronic systems. Lattice QCD can also predict the properties of strongly interacting matter in thermodynamic equilibrium at high temperatures and densities such as the transition temperature and equation of state (EoS). These calculations are relevant for the relativistic heavy ion experiments at RHIC and LHC, and will be discussed below as "Hot QCD". While cold and hot lattice QCD calculations address somewhat different physics objectives, they are technically similar. Furthermore, the synergy between the lattice QCD calculations relevant for the US NP Program and the lattice calculations relevant for HEP program is fully exploited. In 1999, the USQCD consortium was created for this purpose. The development of the software and operation of the dedicated computing hardware, as well as the allocation of the external resources, proceeds through the USQCD consortium. In the past, the software development for lattice QCD was funded through joint HEP-NP SciDAC-2 project. Presently, the software development of nuclear physics lattice QCD is funded by the SciDAC-3 grant "Computing properties of hadrons, nuclei and nuclear matter from QCD", although some of the software development activities are coordinated between HEP and NP through USQCD.

In addition to lattice QCD simulations of stationary nuclear matter properties, the quantitative description of the highly dynamical, rapidly expanding matter created in relativistic heavy-ion collisions requires computationally intensive simulations of a different kind that address the non-equilibrium and near-equilibrium evolution of dense systems of quarks and gluons. These will also be discussed in the "Hot QCD" subsection.

2.1.1 Cold QCD

The main focuses of lattice QCD for cold nuclear physics are the calculations of the bound state spectrum of QCD, the study of the internal structure of the nucleon and nuclei, the study of the nuclear forces, including two and higher body interactions, and the role of QCD in tests of fundamental symmetries and the Standard Model. The calculations of the QCD spectrum are directly relevant for the interpretation of the experimental results from the GlueX experiment as well as for the search for so-called missing resonances at CLAS experiments. Lattice calculations of nucleon structure are closely related to the experimental efforts in the RHIC spin physics program, measurements of generalized parton distributions at JLab and measurements of single spin asymmetries in the COMPASS experiment. Additionally, precise calculations of nucleon and nuclear axial form-factors will allow future measurements of neutrino oscillation parameters, and the possible discovery of new neutrino states at the Long Baseline Neutrino Facility (LBNF). Calculations of the scalar and tensor charges of nucleons and nuclei, the nucleon EDM (arising both from the QCD θ-term and from higher dimension operators) and novel current matrix elements are vital input for the experimental program in fundamental symmetries (neutron β decay, nucleon, nuclear and atomic EDMs, $\mu \rightarrow e$ conversion and proton decay experiments), and for dark matter direct detection. Improved understanding of the nuclear forces, for example the three-neutron forces and hyperon-nucleon interactions, is important for ab initio nuclear structure calculations and the FRIB experimental program, and also experiments at the J-PARC and FAIR facilities. The role of parity violation in these forces is also central to the $np \rightarrow d\gamma$ experiment at the Spallation Neutron Source (SNS). Lattice QCD is a way to, not only understand, but to rigorously refine the interactions between nucleons from first principles. Important developments in the last few years have been the inclusion of quantum electrodynamics in LQCD calculations and the appearance of first results with the physical values of the up, down, strange and charm quark masses. The next generation of LQCD calculations are essential for optimizing the physics impact of the NP experimental program, and for refining the inputs and constraints on nuclear structure and reaction calculations and nuclear astrophysics simulations.

The objective for the next five years is to perform precision Lattice QCD calculations of the properties and structure of the nucleon, of the spectra of mesons (including exotic mesons) and baryons, of the spectra and interactions of the lightest nuclei and hypernuclei, and of the neutron EDM and other inputs needed for the experimental program on fundamental symmetries. The broad range of topics that will be addressed are discussed in more detail in the 2013 USQCD Cold Nuclear Physics and Intensity Frontier Whitepapers. Calculations will be performed at and near the physical quark masses with complete quantification of all associated uncertainties. This will be achieved, in part, through a coordinated effort in gauge-field configuration generation (detailed in the Cold Nuclear Physics Whitepaper) that will provide large, statistically precise ensembles of lattices with a range of masses, volumes and lattice spacings to ensure that all of the important observables are determined with the requisite precision. These will be followed by ensembles that include isospin breaking and electromagnetism. Continued access to Leadership-Class computing resources through the INCITE and ALCC programs and other mechanisms, complemented by capacity computing resources of comparable magnitude, along with support for scientists with expertise in HPC, is critical to achieving these goals.

The major accomplishments of cold lattice QCD calculations since 2007 include:

- Calculations of g_A and other matrix elements such as the scalar and tensor charges (relevant for neutron decay experiments) in the nucleon at, and near, the physical quark masses
- The isovector meson spectrum, hinting towards existence of hybrid mesons
- The spectrum of the lightest nuclei and hypernuclei
- Moments of the generalized parton distributions in the nucleon, and determination of the proton spin decomposition

- Studies of transverse momentum-dependent distributions of the nucleon
- Nucleon-nucleon and hyperon-nucleon phase-shifts
- Determinations of the pion-pion scattering parameters and phase-shifts
- Studies of finite-density Bose-condensed multi-meson systems
- The low lying nucleon and Δ spectra
- Exploratory studies of θ-term contributions to the neutron EDM and of $I = 1$ hadronic parity violation.

2.1.2 Hot QCD

The study of the properties of hot and dense QCD matter is one of the four main areas of nuclear physics research described in the 2007 NSAC Long Range Plan and reaffirmed in the 2013 Long Range Plan Implementation Report. Data taken at the Relativistic Heavy Ion Collider (RHIC) since the year 2000 has shown conclusively that the high temperature phase of QCD matter is a quark-gluon plasma with the characteristics of a "nearly perfect" liquid. This conclusion has been confirmed by measurements in Pb+Pb collisions at much higher energy at the CERN Large Hadron Collider (LHC). The RHIC and LHC experiments have not only shown that the quark-gluon plasma behaves as a nearly inviscid liquid, but also that it is highly opaque to energetic partons, resulting in strong jet quenching.

The physics goal for the next decade is to characterize the properties of this quark-gluon plasma liquid by quantitative extraction of important medium parameters from precision measurements of sensitive observables, including hadron spectra, angular distributions and correlations, jets, and electromagnetic probes. The DOE Performance Measures for High Temperature, High Density Hadronic Matter require: "By 2015, create brief, tiny samples of hot, dense nuclear matter to search for the quark-gluon plasma and characterize its properties." The criterion for the grade "Excellent" demands that "its properties such as temperature history, equation of state, energy and color transport (via jets), and screening (via heavy quarkonium production) are characterized."[1]

A key goal of computational heavy ion theory is to dynamically simulate the space-time development of a heavy-ion collision, from the early time non-equilibrium dynamics, then through a phase that is likely best described by hydrodynamics followed by a freeze-out phase, where the description is switched to a kinetic theory of interacting hadrons. There are, however, several missing ingredients that are needed before a full dynamical simulation can be achieved. A major uncertainty in understanding the physics of heavy ion collisions stems from the description of the very early stages immediately after the collision. Initially the produced matter in a heavy-ion collision is in a state far from thermal equilibrium. Important theoretical questions concern the details of the thermalization mechanisms and the onset of hydrodynamic behavior as well as when and to what extent the formation of a thermalized Quark Gluon Plasma is achieved. If thermalization is indeed achieved at some stage of the heavy ion collisions, the study of the early time dynamics will provide the initial conditions for subsequent hydrodynamic evolution. Furthermore, the properties of the matter can be calculated using lattice QCD. For extracting some properties of strongly interacting matter at quantitative level, e.g. the shear viscosity to entropy density ratio, η/s, a detailed data-to-model comparison is needed. This turns out to be computationally intensive. Below we describe these three computational components of heavy ion physics, commonly referred to as Hot QCD, in more detail.

Thermalization and classical field dynamics: The system created at the very early stages of the collisions is characterized by very large gluon occupation numbers. Therefore, at early times one can use

[1] *Report to NSAC of the Subcommittee on Performance Measures*, August 2008.

classical-statistical field simulations in real time to study the process of thermalization. On a qualitative level, important insight on the thermalization process can already be gained from 3+1 dimensional simulations for the SU(2) gauge group. For realistic calculations one needs to consider the SU(3) gauge group and also include fermions in the simulations. Performing the calculations with SU(3) will need tenfold increase in computing resources. Including fermions in real time simulations could easily increase the computational costs by two orders of magnitude.

Lattice QCD at non-zero temperature and density: For the range of temperatures and net baryon densities that are relevant for the physics of heavy ion collisions, lattice QCD is the only way to provide first principles calculations of the properties of QCD matter near thermal equilibrium. Some basic properties of this matter, such as the chiral transition temperature and the equation of state at net zero baryon density have been recently calculated with controlled uncertainties. Fluctuations of conserved charges are important for understanding the freezeout process in heavy ion collisions, i.e. how to transition from hydrodynamic description to hadronic transport. The next challenge is to perform these calculations at non-vanishing net baryon density which is crucial for the beam energy scan (BES) program at RHIC, and in particular provide some evidence for or against the existence of the critical end-point (CEP) on the QCD phase diagram. For understanding the electro-magnetic, open and hidden heavy flavor probes in heavy ion collisions, lattice QCD studies of meson spectral functions are very important. In essence these encode in-medium properties of different mesons, their dissolution, as well as the rate of thermal photon and dilepton production. Calculations of the transport coefficients in lattice QCD is also of great importance. Such calculations could confirm the experimental findings on the perfect liquid nature of QCD matter in the studied temperature region, and tell us how this perfect liquid emerges from the weakly interacting quark-gluon gas at very high temperatures by studying the temperature dependence of the transport coefficients in a wide temperature range.

Modeling of heavy ion collisions: In recent years the amount of experimental data on heavy ion collisions increased dramatically. To make use of these experimental data to quantitatively characterize the matter produced in such collisions requires a comprehensive data-to-model comparison. It turns out that fluctuations in initial conditions are very important in the description of many quantities, e.g. the anisotropic flow coefficients v_n. It is no longer sufficient to consider only typical initial conditions; rather, one must run hydrodynamic simulations with fluctuating initial conditions on an event-by-event basis. For the late-time evolution, the macroscopic hydrodynamic approach breaks down and has to be replaced by microscopic hadronic transport simulations. These describe the non-equilibrium evolution using stochastic Monte-Carlo sampling techniques that require the hadronic transport codes to be run many times to obtain statistically reliable results. To reduce computational cost, in the past (2+1)-dimensional hydrodynamic calculations that assume boost invariance along the beam direction have been used. These do not allow study of the effects of longitudinal fluctuations on the final state observables. At lower center of mass energies, as explored in the RHIC Beam Energy Scan (BES) program, the assumption of boost invariance is no longer valid. To realistically describe heavy-ion collisions at such lower collisions, and to explore the consequences of density and momentum fluctuations along the longitudinal direction, (3+1)-dimensional hydrodynamic calculations coupled to microscopic hadronic transport will be necessary. This will increase the computational needs by at least a factor of ten. Simulations of the realistic medium evolution are also an important ingredient for the study of high momentum and electromagnetic probes. These studies require significant additional computing resources because calculations of some observables, such as reconstructed jets, require the use of complex Monte-Carlo methods.

To conclude this section, we briefly summarizing the major accomplishments of the computational component of Hot QCD. These include

- The detailed understanding of the chiral and deconfinement aspects of the quark-hadron crossover in QCD at small net baryon density and the determination of the chiral crossover temperature, $T_c =$

$154(9)$MeV

- The definitive determination of the equation of state at zero net baryon density for physical quark masses and in the continuum limit
- A study of fluctuations of conserved charges on coarse lattices that allows to explore the freeze-out conditions in relativistic heavy-ion collisions
- Exploratory calculations of the hadronic spectral functions and determination of some quark-gluon plasma transport coefficients, neglecting the effect of dynamical quarks
- Significant progress in understanding the early time dynamics in heavy-ion collisions in the framework of gluon saturation, including the role of turbulence in the process of thermalization
- (3+1)-dimensional relativistic viscous hydrodynamic simulations of heavy-ion collisions using realistic initial conditions, and the quantitative extraction of the shear viscosity to the entropy density ratio, η/s.

2.1.3 Computational Resources

A significant fraction of the computational resources for lattice QCD come from the dedicated capacity hardware operated by USQCD and funded jointly by NP and HEP. The total USQCD resources amount to 468M CPU core-hours/yr and 8.9M GPU hours/yr. For typical lattice QCD applications the amount of the GPU time corresponds to about 265M CPU core-hours/yr. Roughly half of the total USQCD resources are available for nuclear physics. Therefore 366M core-hours/yr are available for nuclear physics calculations from USQCD hardware. Cold and Hot lattice QCD researchers also receive considerable resources from INCITE, NERSC, NSF, as well as from institutional computational centers. In 2014, LQCD was the largest INCITE award in 2014 (340M core-hours); the NP part was 100M core-hours. The total amount of computer time available for cold and hot lattice QCD in 2014 is approximately 600M core-hours. The longer term computational needs of Cold and Hot lattice QCD are summarzied in the USQCD whitepapers: Lattice QCD for Cold Nuclear Physics and Computational Challenges in QCD Thermodynamics.

The computational resources required to study the early time dynamics and for modeling heavy-ion collisions are considerably smaller than those for lattice QCD. However, since these simulations are presently not highly parallelized and thus do not efficiently exploit HPC resources, most leadership class facilities are closed to users performing them. As the various types of dynamical evolution algorithms have recently reached a stage of maturity, the need for capacity computing resources for heavy-ion collision modeling will dramatically increase in the future. The computational requirements of studying thermalization and early time dynamics of purely gluonic systems with SU(3) interactions require about 2M core hours. Including fermions may increase the computational costs by a factor of 100, thus requiring 200M core hours. Future studies of early time dynamics will therefore need to involve large scale parallelization and require allocation of leadership class resources. – Realistic modeling of the space-time evolution of heavy-ion collisions is a very challenging task. Modeling a minimal set of collisions using a hybrid code coupling (3+1)-dimensional hydrodynamics of the quark-gluon plasma stage with a microscopic transport model for the hadronic stage consumes in excess of 5M core hours. For a comprehensive model/data comparison, and the quantitative extraction of the key medium properties together with the theoretically poorly constrained spectrum of initial-state quantum fluctuations from the experimental data, the required computational resources will increase by factor of 10. Such calculations are not amenable to massive parallelization but can be performed on medium-size capacity-computing resources. However, access to such resources needs to be significantly improved over the next few years.

2.1.4 Personnel, Funding, Leveraging

Investment in the personnel responsible for software development will be critical to ensure a viable lattice QCD program, and re-assure leadership-class facilities that computational resources provided to the community will be optimally exploited (thereby leading to at least Moore's Law growth in resources). If the resources available to this component of nuclear physics research increase no faster than Moore's law, the community will likely be unable to synchronously address the questions relevant to the experimental programs. Further significant algorithmic improvements will be required to reach this goal. At the same time, optimizing codes for the new hardware architectures will be increasingly difficult. This will require at least the same level of support that was provided by SciDAC-2, and hence a significant increase in the level of support currently provided by SciDAC-3.

2.2 Nuclear Structure and Reactions

Large-scale nuclear physics computations dramatically increase our understanding of nuclear structure and reactions and the properties of nucleonic matter. The physics research ranges from studies of the nuclear interaction to critical processes in light and heavy nuclei and nuclear astrophysics. In light nuclei, the focus is on *ab initio* calculations of reactions, including fusion and electroweak processes. *Ab initio* techniques are used to systematically reduce uncertainties and make predictions for experimentally inaccessible data and processes, including reactions on unstable nuclei. Computational approaches are being expanded to reactions on medium-mass nuclei. In heavy nuclei, the focus is on structure and reactions using density functional theory and its extensions. The density functional work is closely tied to *ab initio* studies of experimentally inaccessible systems such as neutron drops to enhance the predictive capabilities. Important areas of research are the structure and decays of very neutron-rich nuclei, the dynamics of the fission process in heavy nuclei, and the structure of neutron star crusts. In the area of fundamental physics, the focus is on double-beta decay, hadronic weak interaction studies, and neutrino scattering from nuclei.

The current Quantum Monte Carlo, Configuration Interaction, Coupled Cluster, and Density Functional codes – used in advanced nuclear structure computations – now scale efficiently to the largest computers available. Computational studies of the strongly correlated matter found in nuclei and neutron stars impact experimental programs throughout the US, including FRIB (structure of neutron-rich and proton-rich nuclei and related astrophysical environments; fundamental interaction studies with rare isotopes), ATLAS and other low-energy nuclear physics facilities (structure and reactions of nuclei and nuclear astrophysics), TJNAF (neutron distributions in nuclei, few body systems, and electroweak processes), NIF (light-ion thermonuclear reactions in a terrestrially controlled plasma environment), MAJORANA, EXO, and FNPB (neutrinoless double-beta decay and physics beyond the Standard Model), LANSCE (studies on the properties of fission), and other nuclear physics and astrophysics facilities. Accurate solutions of the strongly interacting quantum many-body systems will yield new insights and the ability to calculate phenomena, processes, and states of matter that are difficult or impossible to measure experimentally, such as the crust of a neutron star or the core of a fission reactor.

Computational nuclear structure is a key ingredient in sensitive experimental tests of fundamental symmetries and constraints on the Standard Model and physics beyond it. A search for neutrinoless double beta decay at a sensitivity corresponding to the inverted mass hierarchy is one of the flagship goals on the NP program on symmetries. Both the interpretation of such experiments and the design (e.g., required target mass) depend crucially on nuclear matrix elements. Neutrino interactions with nuclei play a critical role in astrophysical environments such as supernovae, and also in terrestrial measurements in reactor and accelerator neutrino experiments. In many cases our understanding of the underlying neutrino cross sections are a limiting factor in the present and future experimental program. Improved large-scale nuclear structure calculations are of great importance to the field, including studies of the impact of truncations in various approaches, and of the role of two-nucleon currents in the various approaches. Currently the most stringent limit on EDMs for a large class of theories come from measurements using complex nuclei, such as ^{199}Hg. The time-reversal violation is dominated by nuclear polarization arising from CP-odd NN interactions. Thus calculations of both the NN interaction, using LQCD as discussed above, and of nuclear Schiff moments, using many-body nuclear theory, are need to connect the observations with fundamental theory. Atomic trapping experiments at FRIB could extend atomic EDM measurements to systems where polarizabilities could be enhanced by 4–6 orders of magnitude by special degeneracies and other effects. Exploiting such cases is an important opportunity for NP. In addition to these priorities, nuclear structure is very important in experimental searches for dark matter and nuclear beta decay experiments probing physics beyond the standard model, either searching for novel interactions, testing unitarity of the CKM matrix or searching for non-standard asymmetries. In all of these cases computational nuclear physics will play important role in enhancing the reach of the NP experimental program.

Computational nuclear structure and reactions in the U.S. has advanced significantly through the UNEDF SciDAC-2 project, which joined the forces of nuclear theorists, computer scientists and applied mathematicians. Integral to the UNEDF project was the verification of methods and codes, the estimation of uncertainties, and assessment of validity. Methods to verify and validate results included the cross checking of different theoretical methods and codes, the use of multiple DFT solvers with benchmarking, and the confrontation of *ab initio* functionals with *ab initio* structure using the same Hamiltonian. The UNEDF project helped form a coherent nuclear theory community, opened up new capabilities, fostered transformative science resulting in high-visibility publications, and advanced the careers of many junior scientists. The UNEDF experience has been a springboard for advancement, as UNEDF postdocs have obtained permanent positions at universities and national labs. Another new aspect to the low-energy theory effort driven by SciDAC is the greatly enhanced degree of quality control. The successor of UNEDF, the NUCLEI SciDAC-3 project, bridges the scales from hadronic interactions to the structure and dynamics of heavy nuclei, to neutron stars within a coherent framework. NUCLEI is strongly coupled to all of the SciDAC Institutes; aided by these connections, the collaboration develops novel computational tools that are specifically needed to accomplish the physics goals.

The major scientific achievements in computational nuclear structure and reactions since 2007 include:

- Dramatic advances in *ab initio* studies of light-ion fusion reactions, combining the no-core shell model and resonating group techniques.
- The first *ab initio* calculations of important and unique light nuclear states and transitions, including the Hoyle state of ^{12}C and the life-time of ^{14}C.
- Investigations of neutral weak current two-body contributions in inclusive scattering from ^{12}C.
- Critical accomplishments in *ab initio* calculations of medium-mass neutron-rich nuclei emphasizing the importance of three-nucleon interactions and particle continuum.
- Improved density functionals (UNEDF0, UNEDF1, and UNEDF2) that simultaneously reproduce *ab initio* calculations of inhomogeneous neutron matter, and significantly advance our abilities to reproduce fission barriers and half-lifes.
- Calculations, including error estimates, of the full range of atomic nuclei, based upon these improved density functionals.
- Calculations of the neutron matter equations of state, relevant to physics of neutron-rich nuclei and neutron star crust, in both *ab initio* and density functional calculations.
- Important impacts on other fields, particularly cold atom physics, in both *ab initio* and density functional calculations.

2.2.1 Computational Resources

These advances were made possible by a rapid increase in our ability to use the largest-scale computational resources, and in available computational resources. Large-scale usage in nuclear structure and reactions rose from 80M core-hr/year in 2009 to 345M core-hr/year in 2014, with over 500M core-year hour projected for 2015. Nuclear structure and reactions researchers receive considerable resources from INCITE, NERSC, NSF (TeraGrid, NICS), as well as institutional computational centers. The INCITE resources, particularly those at OLCCF, have increased rapidly. In 2014, our nuclear theory INCITE was fourth largest out of 59 awards for 2014.

In addition, there is a tremendous need for additional computational resources, as indicated by the expected 2015 requests. In the future, our ability to use the largest-scale computers will require additional investments in manpower, particularly as we transition to new architectures and work to maintain a portable code-base across emerging architectures.

2.2.2 Personnel, Funding, Leveraging

The issue of personnel in computational nuclear structure and reactions is critical. UNEDF brought together a big fraction of the nuclear theory community, and represented a significant increase in support for junior scientists. This was possible only through a tremendous leveraging of the base programs in Nuclear Physics and ASCR, and strong support from NNSA. The UNEDF SciDAC-2 project typically involved approximately 50 scientists in physics, computer, and applied mathematics, including full support for 11 students and 19 postdocs per year. Senior scientists were typically supported only a modest amount or not at all, depending upon particular laboratory and university situations. The full budget for UNEDF was $3 M per year, with NP, ASCR, and NNSA contributing $ 1 M each. Other nuclear structure and reactions collaborations, utilizing significant computing resources, include NuN, TORUS, and SA-NCSM.

Due to the limited resources, the NUCLEI proposal was scaled back to $ 2.7 M per year (NNSA contribution is $ 0.5 M), and another scientifically strong SCIDAC-3 proposal in nuclear structure and reactions was not funded. The current NUCLEI funding is $ 2.775 M per year, including recent supplemental funding of $ 375 K for $0\nu\beta\beta$ calculations. The budgetary reduction, combined with inflation, will reduce our ability to support the next generation, endangering the program in future years. Currently, NUCLEI support approximately 5 students and 13 postdocs per year; several additional students and postdocs are supported at a partial level. Present funding is clearly insufficient to ensure a strong future for the field, support from other sources, including NP and ASCR base funding, and NNSA, will have to provide the rest.

We estimate a need for an additional 10 positions over 5 years in nuclear structure /reactions and associated computer science and applied math. Approximately 1/2 of these positions would be nuclear physicists and 1/2 computer science and applied math scientists. These new positions are critical to enable the effort to scale to the largest-scale machines heading toward exascale, and to support the FRIB and related experimental program in a timely manner. Given sufficient support, we can fundamentally change the future of the field, as demonstrated by the recent history outlined below.

2.3 Nuclear Astrophysics

A primary goal of nuclear astrophysics, and of Nuclear Physics at the DOE, is the explanation in detail of the origin of the elements and their isotopes. Overwhelmingly, the elements of Nature are produced in stars, either during quiescent thermonuclear burning stages or explosively, as in core-collapse and thermonuclear supernovae, gamma-ray bursts, X-ray bursts, and novae. Most of the elements up to the iron peak are ejected in supernova explosions. The r-process has been suggested to occur in core-collapse supernovae (or perhaps some related death throe of a massive star), but heavy elements may be ejected when neutron stars in binaries merge forming gamma-ray bursts. The proton-rich nuclei of the rp-process seem to be produced in X-ray bursts, although processes operating in supernovae may also produce proton-rich species.

Whatever the details of their origin and history, an understanding of stellar explosions and stellar evolution is central to an understanding of the abundance pattern of nuclei around us. Stellar evolution calculations for the entire range of stellar masses involve nuclear reaction rates, generated theoretically or by nuclear experiment, and the complexity of three-dimensional turbulence and magnetic interactions. Stellar explosions are always multi-dimensional, requiring state-of-the-art radiation/hydrodynamical simulations with significant nuclear physics input. As a consequence, much of nuclear astrophysics entails sophisticated numerical simulations employing the latest computational tools and the most powerful supercomputers of the DOE complex to address key goals of the Office of Nuclear Physics.

Computational nuclear astrophysics provides essential support for important components of the DOE Office of Nuclear Physics experimental program. The astrophysics of neutron-rich nuclei is one of four scientific "legs" of the Facility for Rare Isotope Beams. FRIB experiments will determine the masses and beta-decay rates along much of the r-process path. However, the exact path of the r-process, and which nuclear data is therefore of greatest interest, depends strongly on the details of the astrophysical environment (or environments) where these isotopes are forged. Since core-collapse supernovae (CCSN) and merging neutron stars are leading candidate sites for the r-process, the increased fidelity of CCSN and merging neutron-star simulations can be viewed as an indispensable complement of FRIB efforts to better define the nuclear physics of the r-process. Simulations of X-ray bursts fulfill a similar function for experiments exploring the rp-process along the proton drip line, identifying the isotopes created when these explosions occur on the surface of neutron stars.

The flagship experiment in the DOE's Intensity Frontier program is the LBNF mega-detector that will allow experimentalists to follow the neutrino light curve from the next galactic supernova. There are novel neutrino flavor phenomena associated with supernovae – collective flavor oscillations sensitive to the neutrino hierarchy and a second MSW crossing associated with the mixing angle θ_{13} – that could emerge as new physics from detailed comparisons of supernova theory with observation. It is thus of great importance to develop models of supernovae that can simultaneously treat the explosion mechanism and the underlying neutrino microphysics realistically, to assess the effects on nucleosynthesis and to guide the interpretation of neutrino signals that will be seen in underground detectors.

The nuclear equation of state is a third intersection with the DOE experimental program. A JLab program to constrain the nuclear symmetry energy, and, thus, the EoS for neutron-rich matter, by measuring the weak charge radius of Pb has yielded first results. While the JLab experiment tests the symmetry energy at about half nuclear density, this data point still provides an important experimental test of the nuclear theory used in modeling the high-density supernova core and neutron stars.

To accomplish the computational astrophysics goals and projects articulated in the 2007 DOE Long-Range Plan for Nuclear Physics, the tight collaboration of computer scientists, applied mathematicians, and nuclear physicists to create efficient, highly-scalable, parallel simulation capabilities is required. Furthermore, access to resources on the petascale (in the near term), then exascale (on the decadal horizon), will be necessary to fully realize the stated astrophysical science goals of the Office of Nuclear Physics.

2.3.1 Accomplishments

Major scientific milestones since 2007 include:

- The development of 3D codes with spectral neutrino transport for the study of core-collapse supernovae, superseding the previous 2D capability
- The emergence of realistic whole-star 3D models for thermonuclear supernova explosions (Type Ias) with the deflagration to detonation transition

- The development of a new low Mach number algorithm for modeling the convective burning phases preceding the explosion in many of astrophysical events
- The publication of the first detailed 2D and 3D stellar evolution calculations, going beyond the traditional 1D models with ad hoc mixing-length convection theory
- The first neutron-star/neutron-star merger simulations in full general relativity with magnetic fields.

2.3.2 Computational Resources

Many of these advances were made possible by generous access to the largest-scale computational resources available via DoE Office of Science and NSF. In 2014, these include 135M CPU core-hours from the INCITE program and more than 58M CPU core-hours at NERSC. NSF's PRAC program provided 181M CPU core-hours to projects of interest to nuclear astrophysics and the NSF XSEDE program added more than 30M CPU core-hours. Thus the total for nuclear astrophysics amounted to more than 400M CPU core-hours in 2014.

However, there is a growing demand for additional computational resources as three-dimensional radiation-hydrodynamic and massive nuclear network simulations for a variety of key nuclear astrophysics problems become conceivable and feasible. Significant investments, in both scientific and applied-mathematics manpower, are needed to maintain this momentum. These investments are all the more crucial as we transition to new computational architectures and programming models beyond the petascale, yet nuclear astrophysics was shut out of most recent round of the SciDAC program. Steps must be taken to redress this shortfall if nuclear astrophysics is to remain capable of utilizing the coming generations of high-performance computers.

2.3.3 Personnel, Funding, Leveraging

As stated above, it is critical that the personnel issues needed to realize the potential of this next generation of nuclear astrophysics capabilities be addressed. In this regard, it is disappointing that no astrophysics proposals were funded during this last (2013) SciDAC round. This has left numerous groups scrambling to find alternate funding for their junior team members and to maintain their scientific viability. Currently, it is estimated that there are $\sim$50 mid-level researchers in the various previous SciDAC teams who are affected (directly or indirectly). While most of them will find (or have found) alternate employment and leveraging by alternate sources was indeed necessary in the past, the effect of this progressive defunding by the DOE over the last years has severely compromised the United States' ability to realize the progress envisioned in the context of previous DOE investments. We estimate that, factoring in University, NSF, and National Laboratory leveraging, nuclear astrophysics will now need additional support for $\sim 10-20$ junior nuclear astrophysicists and computational scientists to reestablish the momentum previously enabled by past DOE support. These new positions are crucial as we head toward the exascale, and to support the nuclear astrophysics component of the Office of Nuclear Physics experimental program. Failure to address this critical need will, over time, cede leadership in computational nuclear astrophysics to Japan and Germany, both of which are ramping up investment to establish both 1) the critical masses in manpower and 2) the access to the heaviest iron needed to realize the transformative potential exascale presents.

Chapter 3

Workforce, Education, Training and Awards

3.1 Computational workforce: SciDAC and complementary programs

The U.S. DOE SciDAC program has enabled significant advances in computational nuclear physics within the U.S. in the past decade, precisely because it has brought together leading domain specialists (nuclear physicists) and computational scientists to focus on algorithmic improvement as well as improving domain codes for NP. The SciDAC programs enable the support of a very special class of collaborators - ones not typically covered under base funding. Notably, the algorithmic advances and highly optimized codes that SciDAC has facilitated, have been key to the leveraging of substantial leadership-class computational resources.

The SciDAC-funded positions that have been enabled are cost effective and highly leveraged as the many domain scientists with which they interact are supported out of NP base funding. For example, the UNEDF SciDAC-2 and NUCLEI SciDAC-3 projects in nuclear structure and reactions brought together a large fraction of the nuclear theory community, and represented a significant increase in support for junior scientists (full support for 11 students and 19 postdocs per year in UNEDF). But this was possible only through a tremendous leveraging of the base programs in Nuclear Physics and ASCR, and strong support from NNSA.

In FY13, the historical level of NP funding for SciDAC was reduced by a factor of nearly three; *key areas of computational nuclear physics now lack support for the development of the computational infrastructure, with the potential loss of leadership not only in the computational domain, but also in the corresponding experimental and theoretical domains.* In addition, scientifically strong SCIDAC-3 proposals in astrophysics and in nuclear structure and reactions were not funded. Present funding is insufficient to ensure a strong future for the field and it is critical to reverse this striking decline.

The SciDAC program provides the bulk of support at present for software and algorithmic support of computational nuclear physics, but there are needs not covered by SciDAC. For example, there are other nuclear structure and reactions collaborations utilizing significant computing resources, including NuN, TORUS, and SA-NCSM that are not directly supported by SciDAC. There are also other programs that provide, or could potentially provide, additional contributions such as the NSF Software Infrastructure for Sustained Innovation programs, Intel Parallel Computing Centers, and other partnerships with industry. In QCD in particular, the strong partnership with nVidia has led to very effective use of GPUs (former LQCD researcher Mike Clark was hired by nVidia and tasked with supporting LQCD and related fields). In addition, the various DOE computing facilities have programs for postdoctoral positions for candidates working in scientific areas aligned with leadership-class computing. One recent example is the NERSC

Exascale Science Applications Program (NESAP) that will fund up to eight postdocs in various scientific domains. There are and have been similar opportunities at other labs and computational nuclear theory PhDs have proven to be very competitive when seeking these domain-science-related postdoc positions. These smaller scale programs are very important and strongly encouraged as they allow more dynamic response to the emerging needs of the field.

Computational training in nuclear physics can also help meet the workforce needs of industry, where high-level computational skills are increasingly in demand. An example of how to facilitate the bridge between computational nuclear physics and industry is Insight Data Science Fellows Program, which is an intensive six week post-doctoral training fellowship for the transition from academia to data science. The success rate of this program is spectacular and fellows include a number of computational NP researchers (e.g., Mike Kuhlen, UC Berkeley).

3.2 Education and training initiatives

The education of future computational scientists is of major importance in the field as these junior staff will be the future leaders. Recognising this, the various subfields of computational nuclear physics have developed schools for graduate students and postdocs or facilitated participation in training programs with a more general focus. A strong feature of the schools is the international focus, with many US schools having significant numbers of foreign participants and many US students participating in training abroad. Here we give examples of some of these recent activities.

Summer schools focused on lattice QCD, which help train the next generation of researchers in the field, include

- INT Summer School on Lattice QCD for Nuclear Physics, August 6 - 24, 2012, INT, Seattle.
- Lattice QCD, Hadron Structure and Hadronic Matter, Dubna International Advanced School of Theoretical Physics-Helmholtz International School, September 5 - 17, 2011, Dubna, Russia
- Summer School on Lattice QCD and its Applications, August 8 - 28, 2007, INT Seattle

Lattice QCD researchers from the US were actively involved as organizers and/or lecturers in each of these schools.

The TALENT initiative (Training in Advanced Low Energy Nuclear Theory) aims at providing an advanced and comprehensive training to graduate students and junior researchers in low-energy nuclear theory. This initiative aims at developing a broad curriculum that will provide the platform for a cutting-edge theory to understand nuclei, nuclear reactions, and nuclear astrophysics, all including hands-on training in computational aspects. Six TALENT schools have been offered through summer 2014, including one devoted solely to Computational Many-body Methods for Nuclear Physics (at the ECT* in summer, 2012), with three more planned for 2015. The educational material generated for these schools are being collected in the form of WEB-based courses, textbooks, and a variety of modern educational resources. Other TALENT activities include the UiO-MSU-ORNL-UT School on the computational quantum many-body problem. The topics included: computational quantum chemistry, dynamic multi-threading in numerical computing, lattice simulations, advanced parallelization for multi-core and GPU processing for linear systems and eigenvalue calculations, and challenges and approaches for heterogeneous HPC.

Nuclear theory students also take advantage of university courses on high performance computing offered at several universities, including FSU, LSU, OSU, and UTK. A welcome development is that universities are now also embracing the concept of a specialization in computational/data-science in which a student is co-supervised at some level by a high-performance computing expert in addition to their domain science advisor. A successful example is the UC Berkeley Designated Emphasis in Computational Science and Engineering.

An important tool for education and recruitment of the most outstanding workforce are the various prestigious computational science focused fellowships for graduate students such as the DOE Computational Science Graduate Fellow (Krell) and the NNSA SSGF Fellowship. These are enormously beneficial fellowships that expose students to computational problems at a deep level and involve a significant research stay at a national laboratory. A number of current graduate students in the field currently hold such fellowships (Scott Moreland (Duke) and Adam Richie-Halford (Washington)). We strongly encourage expansion of such programs.

One possibility to further enhance computational science within NP would be to develop a competitive computational NP postdoc program, perhaps modeled along the lines of the LHC Theory Initiative Fellowships. In such a program, applicants would partner with a research group and propose to address a particular computational problem. These proposals would then be competitively reviewed by experts from the field with the winner being offered a prestigious postdoctoral position. Such a program could address the software and algorithmic demands of areas with smaller needs than the full-fledged SciDAC programs.

3.3 Career paths of junior researchers

Since the last long range planning exercise, many junior researchers with a strong emphasis on computational nuclear physics have obtained tenured/tenure track positions at universities and laboratories in the US and abroad. The JLab bridge and joint positions, and the RIKEN-BNL bridge positions, have played an important role in facilitating these hires particularly in QCD and heavy ion physics. The future FRIB theory center will provide similar stimulus in the nuclear structure and astrophysics communities. Recent junior faculty hires and permanent laboratory positions in the field are listed in Tables 3.1 and 3.2. It is critical to maintain this stream of junior scientists, and to provide a balance across types and fields of positions. In addition, numerous students and postdocs in computational nuclear physics have gone on successfully to positions in industry, where their skills are in high demand (so much so that national labs have had difficulties in making competitive offers [1]).

Table 3.1: Junior faculty (including Research Assistant Professor, RAP), permanent lab staff, and 5 year positions from 2012-present with strong computational emphasis or significant support from a SciDAC project during postdoctoral position(s). The fields are given as Cold QCD (mostly lattice), Hot QCD that includes lattice and relativistic heavy-ion phenomenology, nuclear structure and reactions (NS/R) and computer science (CS).

Person	Field	Institution
2014		
Hasan Metin Aktulga	(CS)	Michigan State U
Heiko Hergert	(NS/R)	Michigan State U, 5 year
Nobuo Hinohara	(NS/R)	U of Tsukuba
Jason Holt	(NS/R)	TRIUMF
Jeremy Holt	(NS/R)	U Washington, 5 year
Alessandro Lovato	(NS/R)	Argonne
Stefan Meinel	(Cold QCD)	Arizona, RIKEN
Ken Nollet	(ASTRO/NS/R)	SDSU (offered)
Hiroshi Ohno	(Hot QCD)	U of Tsukuba
Bjoern Schenke	(Hot QCD)	Brookhaven
Andrew Steiner	(ASTRO)	U of Tennessee
Christopher Thomas	(Cold QCD)	U of Cambridge
2013		
Silas Beane	(Cold QCD)	U of Washington
Michael Forbes	(NS/R)	Wash State U
Ming Gong	(Cold QCD)	Institute of High Energy Physics (China)
Elena Litvinova	(NS/R)	Western Michigan U
Tom Luu	(Cold QCD)	Jülich
Hannah Petersen	(Hot QCD)	Frankfurt
Andre Walker-Loud	(Cold QCD)	William & Mary/ JLab
Frank Winter	(Cold QCD)	JLab
2012		
William Detmold	(Cold QCD)	MIT
Heng Tong Ding	(Hot QCD)	Wuhan U
Alex Gezerlis	(NS/R)	Guelph
Harvey Meyer	(Cold and Hot QCD)	University of Mainz
Junchen Pei	(NS/R)	Pekin U
Masha Sosonkina	(NS/R)	Old Dominion University
Michael Strickland	(Hot QCD)	Kent State
Joseph Wasem	(Cold QCD)	LLNL

Table 3.2: Junior faculty (including Research Assistant Professor, RAP), permanent lab staff, and 5 year positions from 2008-2011 with strong computational emphasis or significant support from a SciDAC project during postdoctoral position(s). The fields are given as Cold QCD (mostly lattice), Hot QCD that includes lattice and relativistic heavy-ion phenomenology , nuclear structure and reactions (NS/R) and computer science (CS).

Person	Field	Institution
	2011	
Takumi Doi	(Cold QCD)	RIKEN
Joaquin Drut	(NS/R)	UNC
Stefano Gandolfi	(NS/R)	LANL
Pieter Maris	(NS/R)	Iowa, RAP
Markus Kortelainen	(NS/R)	U Jyväskylä, RAP
Gustavo Nobre	(NS/R)	NNDC, Brookhaven
Mark Paris	(Cold QCD)	LANL
Brian Tiburzi	(Cold QCD)	CCNY, RIKEN
Paul Romatschke	(Hot QCD)	Colorado
Ionel Stetcu	(NS/R)	LANL
Jun Terasaki	(NS/R)	U of Tsukuba
	2010	
Andrei Alexandru	(Cold QCD)	GWU
Shinji Ejiri	(Hot QCD)	Niigata University
Swagato Mukherjee	(Hot QCD)	BNL
Claudio Pica	(Hot QCD)	University of Southern Denmark
Nicolas Schunck	(NS/R)	LLNL
Stefan Wild	(CS)	ANL
	2009	
Sonia Bacca	(NS/R)	TRIUMF
Michael Forbes	(NS/R)	INT 5 year
Huey-wen Lin	(Cold QCD)	U of Washington 5 year
	2008	
Saumen Datta	(Hot QCD)	Tata Institute
William Detmold	(Cold QCD)	William & Mary/ JLab
Gautam Rupak	(NS/R)	Mississippi State
Takashi Umeda	(Hot QCD)	Hiroshima U

3.4 Awards and prizes

Researchers in nuclear physics with a strong computational emphasis have received significant recognition for their research accomplishments, both at junior and senior levels. In Table 3.3, we list early career research awards since 2008. Table 3.4 lists some of the major national and international awards garnered by computational nuclear physicists (this list omit a large number of lesser awards and fellowships).

Table 3.3: US early career research awards 2008-present for faculty with strong computational emphasis or significant support from SciDAC during postdoctoral position(s).

Person	Institution	Award	Year
Andre Walker-Loud	William & Mary	DOE Early Career	2014
Bjorn Schenke	Brookhaven	DOE Early Career	2014
William Detmold	MIT	DOE Early Career	2013
Carla Fröhlich	NC State	DOE Early Career	2013
Gaute Hagen	ORNL	DOE Early Career	2013
Paul Romatschke	Colorado	DOE Early Career	2012
Daniel Kasen	UC Berkeley/LBL	DOE Early Career	2012
Andrei Alexandru	GWU	NSF Career	2012
Jo Dudek	Old Dominion	DOE Early Career	2011
Sophia Quaglioni	LLNL	DOE Early Career	2011
Denes Molnar	Purdue	DOE Early Career	2010
William Detmold	William & Mary	DOE NP OJI	2009

Table 3.4: National and international awards 2008-present for faculty with strong computational emphasis or significant support from SciDAC during postdoctoral position(s).

Person	Institution	Award	Year
John Negele	MIT	Herman Feshbach Prize	2014
Bjorn Schenke	Brookhaven	IUPAP prize in Nuclear Physics	2013
Stefan Gandolfi	LANL	IUPAP prize in Nuclear Physics	2013
Péter Petreczky	Brookhaven	Zimanyi Medal	2012
Witold Nazarewicz	Tennessee	Tom Bonner Prize	2012
Steve Pieper and Robert Wiringa	ANL	Tom Bonner Prize	2010
Joaquin Drut	UNC	Kümmel Early Achievement Award	2009
Friedrich K. Thielemann		Hans A. Bethe Prize	2008
W. David Arnett 2009		Hans A. Bethe Prize	2009
George M. Fuller	Chicago	Hans A. Bethe Prize	2013

3.5 Reports on the computational science workforce

Recent reports from ASCAC [1], NSAC [2], and HEPAP [3] assessing workforce development needs in Office of Science research disciplines include documentation and recommendations on workforce needs in computational science. Here we provide some relevant excerpts.

3.5.1 ASCAC

From the Summary of Findings of the ASCAC Workforce Subcommittee Letter (July 23, 2014) [1]:

> "Under the auspices of the Office of Science, a significant fraction of the research and development undertaken at DOE national laboratories is centered on scientific computation. Therefore, maintaining a sufficient workforce in Computing Sciences is critical if the investment in ASCR facilities is to be realized. This report evaluates the Computing Sciences workforce critical for the Office of Science to meet its scientific mission.
>
> Results of data analyzed are that the Computing Sciences workforce recruitment and retention activities are below the level necessary to sustain ASCR facilities and maintain DOEs high standards of excellence for innovative research and development."

Findings and recommendations from Section 8 of the report on existing programs addressing DOE workforce needs says that "The DOE national laboratories face workforce recruitment and retention challenges in all areas of the Computing Sciences that are relevant to their mission . . . " and recommends

> "Preserve and increase investment in the DOE CSGF program while developing new fellowship programs modeled after the CSGF program to increase opportunities for more high-quality students, particularly students from underrepresented populations and demographics, in the computing sciences."

and

> "Establish new fellowship programs, modeled after the CSGF program, for research opportunities in enabling technologies in the computing sciences, including computer science for HPC, large-scale data science, and computational mathematics."

In terms of the ability of existing academic programs to meet workforce demand, the report finds that

> "The successful execution of the DOE mission at present and into the future will require dramatic improvements in multi-disciplinary scientific understanding, where current and future methodologies will need to be coupled with advanced computational modeling and data analytics on state-of-the-art computers. Positions in exascale and extreme computing will create an even greater need for students with a multidisciplinary background that also includes high-performance computing, yet academic programs preparing students in these key areas remains extremely limited."

3.5.2 NSAC

The NSAC report on Assessment of Workforce Development Needs in the Office of Nuclear Physics Research Disciplines (July 18, 2014) [2] gives a detailed overview of needs and issues:

> "Challenges in attracting, training and retaining a talented U.S. workforce in high- performance computing and simulations for nuclear science and its applications. Many of the frontier activities in science, engineering, and technology that address the needs of the nation and industry require state of the art computations and simulations of complex systems on large scale computers. Very few universities have access to the largest computer systems. To be successful in these activities requires a deep knowledge of both the science and the technology that drives the computations and the ability to develop and implement new paradigms in computer science to use effectively the largest computer systems. The Computational Science Graduate Research Fellowships have worked to identify the most talented U.S. students working in high-performance

> computing and the exciting science they want to address. These fellowships require the recipient to spend a practicum at a national laboratory, making these awards truly traineeships. However, this effort falls far short of the needs for these highly talented individuals at the DOE laboratories for fundamental and applied science, that is exacerbated by the highly competitive opportunities in the private sector. Retention is a very serious issue, as scientists well trained in high-performance computing typically have the ability to move into diverse and rewarding areas. Focused research opportunities for undergraduates supported by both NSF and DOE have been successful at recruiting students into computational nuclear physics and nuclear astrophysics, but much more could be done. Extensive training of beginning graduate students could also be very effective in attracting the best and brightest into computational nuclear science. More in-depth training is available through the SciDAC (Scientific Discovery through Advanced Computing) program, which has been extremely successful in training graduate students and postdocs. This program brings together graduate students and postdocs in applied math and computer science and nuclear science to tackle many of the most difficult and important problems. The extended duration of the SciDAC projects allows for in-depth training and collaborations across many universities and national laboratories, to the great advantage of both. The NNSA Advanced Scientific and Computing program similarly trains people across six university centers. Many of the students and postdocs trained through this program have gone on to successful laboratory, university and industrial positions across the U.S. Strong support of these programs is crucial to the future nuclear science workforce. Training in Advanced Low Energy Nuclear Theory (TALENT), endorsed by the FRIB Theory Users Group, hosts modules on high-performance computing and computational tools for nuclear physics."

Specific recommendations relevant for computational nuclear physics are:

> "**We recommend that the Department of Energy and the WDTS create new opportunities for high-performance computational science multi-disciplinary training across SC and in collaboration with NNSA through SciDAC and other initiatives.** Such training is broad since students and postdocs interact with computer scientists and it helps to foster a commitment to computational sciences and collaborations between universities and national laboratories. This training could be realized via 'schools' that bring together nuclear and computer scientists or short training courses hosted by SciDAC or other high-performance computing multi- investigator initiatives."

and

> "**We recommend that WDTS in collaboration with discipline-specific offices such as NP establish prestigious postdoctoral training opportunities in areas of demonstrated need and with an opportunity to couple to DOE laboratories or facilities.** Subfields should include nuclear physics and chemistry; high-performance computational science; accelerator science, engineering and technology. Recipients would be expected to propose a research plan that couples their activities to opportunities at DOE laboratories or facilities. This expectation would serve as an introduction to the broader, multi-disciplinary research at the laboratories and complement their training at their home institution."

3.5.3 HEPAP

The report on HEP Workforce Development Needs (June 30, 2014) [3] notes a workforce deficit in large-scale computing and big data:

“The trend towards many-core, multi-core and GPU chips will have a significant impact on software development for both experiment and theory. These developments are further complicated by the need to deal with multiple levels of parallelism. The development of the physics codes will have to be more closely tied to computing and software developments. Software and algorithm developers will need to be experts to work in this increasingly dynamic, evolving and innovative technology landscape. Traditionally many HEP scientists have worked on the development of physics codes. Many of these scientists will require additional training to keep up with the trends in large-scale computing and big data.

The workforce for large-scale computing and big data in HEP is made up of computing professionals and PhD scientists and domain experts with expert level computing skills. The shortage of workforce trained in these areas has been documented at both FNAL and BNL. FNAL reports that searches for application software developers turn up relatively few applicants, typically 10 to 15, and that most of these are foreign. The turnover rate for computing professionals at Fermilab has averaged 7.4% per year over the last 5 years. An FNAL summary lists this discipline as having the greatest need after problem areas of accelerator science. Likewise, BNL reports that 18 searches in the last three years in the area of Controls/Scientific Computing attracted fewer than ten qualified applicants. Only beamline science had more low-yield searches.

It remains critical that we continue to attract and retain scientists and computing professionals with expert level computing skills to HEP. Some areas where experts are needed include . . . The experts in thesex areas are in high demand, particularly in locations with a thriving high tech economy. The stimulating challenges of the scientific environment have provided sufficient incentive for many of software and computing experts, but the turnover rates in these areas remain high.”

Bibliography

[1] ASCAC Workforce Subcommittee Letter, July 23, 2014, http://science.energy.gov/~/media/ascr/ascac/pdf/charges/ASCAC_Workforce_Letter_Report.pdf

[2] Assessment of Workforce Development Needs in the Office of Nuclear Physics Research Disciplines, July 18, 2014, http://science.energy.gov/~/media/np/nsac/pdf/docs/2014/NSAC_workforce_jul-18-2014.pdf.

[3] HEP Workforce Development Needs, June 30, 2014, http://science.energy.gov/~/media/hep/hepap/pdf/Reports/OHEP_Workforce_Letter_Report.pdf