SciDAC/NUCLEI – Achievements & Plans James P. Vary, Iowa State University

SciDAC-PI Meeting July 22-24, 2015

The Overarching Questions

- How did visible matter come into being and how does it evolve?
- How does subatomic matter organize itself and what phenomena emerge?
- Are the fundamental interactions that are basic to the structure of matter fully understood?
- How can the knowledge and technological progress provided by nuclear physics best be used to benefit society?
 - NRC Decadal Study





Scientific Discovery through Advanced Computing

U.S. DEPARTMENT OF

SciDAC

ENERGY

Office of

Science



The Time Scale

- Protons and neutrons formed 10⁻⁶ to 1 second after Big Bang (13.7 billion years ago)
- H, D, He, Li, Be, B formed 3-20 minutes after Big Bang
- Other elements born over the next 13.7 billion years

Scientific Discovery discogh Advanced Computing

NUclear Computational Low-Energy Initiative



computingnuclei.org (adapted by Gaute Hagen)

Collaborations within NUCLEI



- Multiple direct partnerships formed. This eliminated traditional difficulties of PHY–CS/AM collaborative projects
- Increased collaboration across domains
- Many initially unanticipated blueblue and blue-red links
- Worked well in UNEDF, continuing to work well in NUCLEI
- Collaborations have deepened and expanded
- SciDAC Institutes represented by CS/AM collaborators:
 - FASTMath Ng, Yang
 - SUPER Wild, Sarich
 - o QUEST Higdon, Lawrence

Math /Computer Science

Six case studies to exemplify the SciDAC/NUCLEI achievements and goals Several posters provide greater detail

Highly Parallel Green's Function Monte Carlo (GFMC) for Nuclear Physics

- Physics Problem: Solve for the structure of nuclei: ground state (energy, spin, density profile), spectrum of excited states, magnetic and quadrupole moments, transition rates, ... to compare experiment and make predictions: already theory has predicted new states which have then been observed.
- Computational approach: Use GFMC algorithm to solve the many-nucleon Schrödinger equation for realistic 2- and 3-nucleon interactions
 - Good: extreme accuracy, no basis expansion, no fixed grids
 - Bad: scaling. Need to keep track of every nucleon's spin and isospin states. The number of states is $2A * \begin{pmatrix} A \\ Z \end{pmatrix}$, where A = # nucleons and Z = # protons. Therefore each additional nucleon increases computational load immensely.
- At beginning of SciDAC, GFMC was computing all of the above for A = 10 (beryllium, boron) with a parallel manager/worker algorithm with load balancing (2000 MPI processes on IBM SP at NERSC)
 - The computation depended on doing several (>10) Monte Carlo samples per process.
- Computational challenge:
 - To solve 12 C, an important nucleus, with ~100,000 processes.
 - Needed to use more processes than Monte Carlo samples, so had to break one sample across multiple processes, leading to fine-grain parallelism along with more complex task scheduling and load balancing
 - Wanted to maintain overall manager/worker structure
- Computer Science response via SciDAC:

The Asynchronous Dynamic Load Balancing library



The Asynchronous Dynamic Load Balancing (ADLB) library

- ADLB implements a flexible and scalable scheduling and load-balancing system for work units of varying types, sizes, and priorities.
 - Application processes create and <u>put</u> work units into a shared pool of work units and <u>get</u> and process work units from this pool.
 - There is no manager process; rather, a subset of application processes in constant communication with one another is used to manage the work unit pool in the background.
 - Sophisticated process-load balancing, memory-usage load balancing, and message-traffic load balancing algorithms are used to achieve scalability without burdening the application programmer.
 - ADLB uses MPI and is compatible with MPI usage by the application.
 - Work units can create sub-work units for ADLB to manage
- Physics Results enabled: Solved for ¹²C, including Hoyle state and its decay rate to the ground state, with excellent scaling, providing critical test of the interactions
- Current and future work:
 - More demanding computations: response functions for ¹²C on BG/Q
 - Required: larger work packages => memory management challenges
 - Solution: new distributed memory management library (DMEM)
 - VLMPI: MPI extension for > 2GB messages
 - to be used by DMEM to provide large ADLB work packages
 - Usable by other applications



 $40\square$

See Joe Carlson's poster

Nuclear density functional theory (DFT) is needed to compute properties of atomic nuclei for basic science and applications

Nuclear propertiesFissionFundamental symmetriesNucleosynthesisImage: Strain Strai

- Physics: Develop predictive nuclear DFT to evaluate global properties of nuclei, nuclear fission and for applications to, and beyond, standard model physics and nucleosynthesis
- CS/Math: Develop optimization and uncertainty quantification methods
 - to rigorously estimate and propagate statistical theoretical uncertainties associated with DFT model parameters
 - to assess the information content of experimental observables with respect to current theoretical models
- No UQ framework was available for nuclear DFT because of the high computational cost of DFT calculations

We developed a comprehensive framework to optimize nuclear energy functionals and quantify and propagate the associated statistical uncertainties

- We use our optimization code POUNDerS to determine the coupling constants of nuclear energy density functionals (EDF)
- We compute the posterior distribution of nuclear EDFs by building response functions using Gaussian processes (GPMSA UQ code) and the fast DFT solver **HFBTHO**
- We propagate the resulting statistical uncertainties
 - to provide reliable theory guidance for DOE and NSF experimental programs
 - to identify limitations in predictive power of models.



Propagation of uncertainties



Determination of model parameters

Quantification of uncertainties

See Nicolas Schunck's poster

PRC 82, 024313(2010); PRC 85, 024304 (2012); PRC 89, 054314(2014); J. Phys. G.: Nucl. Part. Phys. 42, 034024 (2015); 034031(2015); 034009 (2015); PRL 114, 122501 (2015)

Need to solve nuclear DFT for complex topologies on multiple scales

Physics problem: *Multiple scales* and *complex topologies* arise in quantum open systems such as triaxial and reflection-asymmetric nuclei, weakly bound halo states, cluster configurations, nuclear fragments produced in heavy-ion fusion reactions, cold Fermi gases, and pasta phases in a neutron star crust. The critical symmetry-breaking physics may be discovered within nuclear density functional theory (DFT).

Computational problem: For the needed full 3D solutions in DFT, calculations for nonspherical geometries are three orders of computational complexity greater than



2D solutions. The desired tool to solve DFT equations in large boxes is an adaptive multi-resolution wavelet analysis *which did not exist in nuclear structure calculations before SciDAC*. Adapting wavelet techniques to the nuclear DFT problem in 3D with numerical stability and with controlled precision presented major challenges that were addressed with the MADNESS suite.

Caption: Illustration of adaptive representations in MADNESS-HFB. (a) The modulus squared of the single-neutron wave function corresponding to the single-particle energy of –5.214 MeV obtained in MADNESS-HF calculations for ¹¹⁰Mo, and (b) the corresponding spectral refinement structure. **Reference**: J.C. Pei et al., Phys. Rev. C. 90, 024317 (2014); Editors' Suggestion

SciDAC/NUCLEI Solution: MADNESS-HFB: adaptive multi-resolution 3D DFT solver

- Introduced an adaptive pseudospectral method for solving self-consistent equations of nuclear density functional theory in 3D, without symmetry restrictions.
- Developed sophisticated object-oriented templates which parse the high-level code into distributed parallel tasks with a multi-thread task queue scheduler for each multi-core node.
- Benchmarked new adaptive multi-resolution solver MADNESS-HFB against 2D coordinate-space solver using B-splines and a 3D solver based on the harmonic-oscillator basis expansion.



- The algorithm is variational and is capable of solving coupled complex-geometric systems of equations adaptively, with functional and boundary constraints, in a finite spatial domain of very large size.
- The new framework solves multi-scale physics challenges in nuclear and atomic problems involving many-particle superfluid systems.

Caption: Single-particle density distributions illustrating (see semilog inset) the superior asymptotic properties of the MADNESS solutions to DFT over conventional techniques. (J.C. Pei et al., Phys. Rev. C. 90, 024317 (2014); *Editors' Suggestion*)



Need *ab initio* solutions of extreme nuclear systems to develop next-generation nuclear energy density functional



Extreme nuclear systems – not accessible in the laboratory

Kepler supernova (NASA)

Nuclei beyond neutron drip line

Neutron star "pasta"

- Physics: Need to solve large neutron drops in external fields to probe the isospindependence of the nucleon-nucleon (NN) and three-nucleon (3N) interactions but these systems exhaust a node's memory for the input M-scheme 3N interactions
- CS/Applied Math: Develop algorithms to compress, store and decompress 3N interactions on Titan's GPUs
- Prior to SciDAC, neutron drops with more than 12 neutrons could not be solved close enough to convergence

Ab initio Extreme Neutron Matter



Physics accomplishments [2]:



Ground state energies of neutron systems in a 10 MeV harmonic oscillator. Solid red diamonds and blue dots signify results with NN+3N interactions derived from chiral effective field theory related to QCD. Inset displays the ratio of NN+3N to NN alone. Chiral results lie between those of AV8'+UIX and AV8'+IL7, two phenomenological interactions [3].



- Demonstrates predictive power of *ab initio* nuclear structure theory.
- Provides results for next generation nuclear energy density functionals
- Leads to improved predictions for astrophysical reactions
- Demonstrates that three-nucleon (3N) interactions in extreme neutron systems are significantly weaker than phenomemology predicts
- Guides future experiments at DOEsponsored rare-isotope facilities

Future work: Extend and apply these techniques to nuclear double beta-decay

46% of authors are grad students

D. Oryspayev, H. Potter, P. Maris, M. Sosonkina, J.P. Vary, S. Binder, A. Calci, J. Langhammer and R. Roth, IEEE 27th IPDPSW & PhD Forum, 1365 (2013)
 H. Potter, S. Fischer, P. Maris, J.P. Vary, S. Binder, A. Calci, J. Langhammer and R.Roth, Phys. Lett. B739, 445 (2014)
 P. Maris, J.P. Vary, S. Gandolfi, J. Carlson, S.C. Pieper, Phys. Rev. C87, 054318 (2013)

Scalable eigensolver required for *ab initio* configuration interaction calculations for nuclear structure

The Challenge in a Nutshell:

Rapid rise in matrix dimensionality translates to increased computational burden



- Physics: Need converged *ab initio* configuration interaction solutions of light nuclei (N_{max} = 10 and above) to describe experiments, predict future observations and validate theoretically-derived strong interactions
- CS/Math: Develop efficient and scalable iterative solvers for extreme-scale eigenvalue problems for nuclear physics (MFDn code)
- Prior to SciDAC, our eigensolvers ran inefficiently on more than 5000 processes

Topology-aware mapping and communication hiding combine for scalable eigensolver in nuclear structure (MFDn)



Topology-aware mapping of processes to the physical processors becomes more important as the gap between computational power and bandwidth widens. Communication groups are optimized through a column-major ordering of processes on the triangular grid [1,2].

- Drastically reduced communication overheads
- Significant speed-ups over earlier version of MFDn (up to 6x on 18,000 cores)
- Almost perfect strong scaling on up to ~200,000 cores
- Future: further reduce communication overheads & make efficient use of more threads/node
- Future: develop/apply complex symmetric matrix eigensolver for coupling to the continuum

See Chou Yang's poster

 H.M. Aktulga, C. Yang, E.G. Ng, P. Maris, J.P. Vary, "Improving the Scalability of a Symmetric Iterative Eigensolver for Multi-core Platforms", CCP&E 25 (2013)
 D. Oryspayev, H.M. Aktulga, M. Sosonkina, P. Maris and J.P. Vary, "Performance Analysis of Distributed Symmetric Sparse Matrix Vector Multiplication Algorithm for Multi-core Architectures," CCP&E (2015) (in press)

Communication Hiding

Flow-chart for multithreaded SpMV computations during the eigensolve phase of MFDn. Expensive communications are overlapped with computations. Explicit communications are carried out over topologyoptimized groups [1,2].



ab initio configuration interaction results enabled by improved eigensolver (MFDn)



Carbon-12 excitation spectra without and with initial chiral 3N interaction for two N_{max} values, compared with experiment [1,3].



Rotational parameters A, a and Ex for ground and excited bands of the Be isotopes [2,3]. Brackets highlight difference between experimental results (horizontal bars) and calculations with extrapolation (parallel triangles) to the infinite matrix limit. Dots represent calculations used for extrapolations.

[1] P. Maris, J.P. Vary, A. Calci, J. Langhammer, S. Binder and R. Roth, "12C properties with evolved chiral three-nucleon interactions," Phys. Rev. C 90, 014314 (2014)
 [2] P. Maris, M.A. Caprio and J.P. Vary, "Emergence of rotational bands in ab-initio no-core configuration interaction calculations of the Be isotopes," Phys. Rev. C 91, 014310 (2015)

[3] J.P. Vary, P. Maris, H. Potter, M.A. Caprio, R. Smith, S. Binder, A. Calci, S. Fischer, J. Langhammer, R. Roth, H.M. Aktulga, E. Ng, C. Yang, D. Oryspayev, M. Sosonkina, E. Saule and U. Catalyurek, "Ab Initio No Core Shell Model – Recent Results and Further Prospects," Proceedings International Conference 'Nuclear Theory in the Supercomputing Era – 2014', Khabarovsk, Russia, June 23-27, 2014. Eds. A.M. Shirokov and A.I. Mazur. Pacific National University (Khabarovsk, Russia, 2015), p. 71.

Nuclear physics challenge: Accurate radii and binding energies for heavier nuclei



Physics Problem: Light nuclei with mass number A≤4, included in the optimization of previous nuclear interactions, exhibit accurate radii and binding energies. Heavier nuclei such as ¹⁶O or ⁴⁰Ca exhibit much too small radii and too large binding energies.

Computational Problem: heavier nuclei with a high computational expense need to be included into the optimization of the nuclear interaction at NNLO.

Computational challenges: Quantify uncertainties and optimize N³LO_{sat}

Solution: simultaneous optimization of NN and 3N forces with input from selected nuclei up to A = 25 (NNLO_{sat}). The computational challenge was met by speeding up coupled-cluster calculations of finite nuclei, and by a scheduler for the computation of a large number of observables required for this study.

[A. Ekström et al., Phys. Rev. C 91, 051301(R) (2015)]



Remaining challenge: Lower the computational expense for more precise coupled-cluster calculations of medium-mass nuclei that enter the objective function. Couple automatic derivatives with optimization routine POUNDerS to understand theoretical uncertainties and correlations between optimization parameters (low-energy constants of the interaction).

Additional PHYS-CS/AM projects underway within SciDAC/NUCLEI:

Prepare for next-generation leadership class facilities:

- NERSC: Maris, Vary, Papadimitriou, Ng, Yang NESAP Award
- OLCF: Hagen, Jensen, et al, CAAR Award
- ALCF: Pieper, Lusk, et al, Proposal for THETA Early Science Award

Integrate POUNDerS with MFDn: Sosonkina, Wild, Maris, Shirokov, et al

Develop workflow software for Coupled Channels applications, Jensen, et al

INCITE = Critical resources for NUCLEI

Ab initia Mathada (CC CEMC NCCM) Soushing

NUCLEI Codes	Application	Production Run Sizes	Resource	Dense Linear Alg.	Sparse Linear Alg.	Monte Carlo	 Density Functional Theory → reasonable time to solution to calculate the entire mass table
	AGFMC: Argonne Green's Function Monte Carlo	262,144 cores @ 10 hrs	Mira			X	2015 3rd largest out of 56
	MFDn: Many Fermion Dynamics - nuclear	260K cores @ 4 hrs 500K cores @ 1.33 hrs	Titan Mira		X		2014 4th largest out of 59
	NUCCOR: Nuclear Coupled-Cluster Oak Ridge, m-scheme & spherical	100K cores @ 5 hrs (1 nucleus, multiple parameters)	Titan		X		2013 2012 6th largest out of 61 6th largest out of 60
	DFT Code Suite: Density Functional Theory, mean-field methods	100K cores @ 10 hrs (entire mass table, fission barriers)	Titan	X			2011 13th largest out of 57 2010 10th largest out of 69
	MADNESS: Schroedinger, Lippman-Schwinger and DFT	40,000 cores @ 12 hrs (extreme asymmetric functions)	Titan	X	x		2009 10th largest out of 66 Largest Project 2008 4th largest out of 55
	NCSM_RGM: Resonating Group Method for scattering	98,304 cores @ 8 hrs	Titan	X	X		0 100 200 300 400 500 Core-hours (millions)

Adapted from slide by Hai Ah Nam, LASL

Conclusions and Outlook

- Disruptive progress in leadership-class machines required new modes of PHYS-CS/AM collaboration to achieve first-rate science
- Case Studies of PHYS-CS/AM collaborations within NUCLEI illustrate successful integrated approach to achieve scientific and technical progress
- Large number of new collaborations formed in NUCLEI and 36 joint PHYS-CS/AM publications produced to date*
- Yearly INCITE awards since 2008 have enabled this scientific progress
- Funding via SciDAC critical for the success of this new paradigm

*see: <u>http://computingnuclei.org</u> for listing of publications