LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-653999

Ab Initio Calculations Of Nuclear Reactions And Exotic Nuclei

S. Quaglioni

May 5, 2014

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

1 Case Study Title: Ab Initio Calculations of Nuclear Reactions and Exotic Nuclei

Prepared by: Sofia Quaglioni, Lawrence Livermore National Laboratory (LLNL)
Contributors: G. Hupin, LLNL; P. Navrátil, TRIUMF; C. Romero-Redondo, TRIUMF; J. Langhammer, Technishe Universität Darmstadt; Robert Roth, Technishe Universität Darmstadt
Refers to the following high-performance computing allocations:
NERSC – Principal Investigator: J. P. Vary, Iowa State University; NERSC Repository: m94
INCITE – "Nuclear Structure and Nuclear Reactions"; Principal Investigator: J. P. Vary, Iowa State University
LLNL Institutional Computing Grand Challenge – "From Nucleons to Nuclei to Fusion Reactions"; Principal Investigator: S. Quaglioni, LLNL

2 Project Description

2.1 Overview and Context

Our ultimate goal is to develop a fundamental theory and efficient computational tools to describe dynamic processes between nuclei and to use such tools toward supporting several DOE milestones by: 1) performing predictive calculations of difficult to measure landmark reactions for nuclear astrophysics, such as those driving the neutrino signature of our sun; 2) improving our understanding of the structure of nuclei near the neutron drip line, which will be the focus of the DOE's Facility for Rare Isotope Beams (FRIB) being constructed at Michigan State University; but also 3) helping to reveal the true nature of the nuclear force. Furthermore, these theoretical developments will support plasma diagnostic efforts at facilities dedicated to the development of terrestrial fusion energy.

Despite nearly 80 years of study, coming to a fundamental understanding of nuclei and their role in the universe remains a central goal of nuclear physics. The origin of this lack of understanding resides in two main challenges characterizing the nuclear many-body problem: i) the interaction between nucleons is complex and is not yet fully understood at a fundamental level; and ii) a fully developed first-principles theory able to provide a unified treatment of a large range of nuclear phenomena (bound-state, scattering, and reaction observables) is still missing.

Our research program is devoted to advancing the state of the art in both these aspects. We apply high-performance computing (HPC) to solve the nuclear many-body problem in terms of constituent protons and neutrons interacting through the most fundamental interactions available nowadays, nucleon-nucleon (NN) and three-nucleon (3N) forces grounded in the fundamental theory of quantum chromodynamics (QCD) within the framework of chiral effective field theory. (The 3N force is a direct consequence of the complexity of the QCD theory.) At the same time, to achieve a robust *ab initio* description of nuclear properties, we are developing a unified and computationally efficient approach to bound and scattering states in light nuclei.

2.2 Objectives for 2017

Until only five years ago, an *ab initio* treatment of light-nucleus fusion reactions was impossible. With our work, we have demonstrated that it is possible to describe complex processes such as the ^{7}Be(p,γ)^{8}B radiative capture[1] (important for its influence on the Standard Solar Model) and the ^{3}H(d,n)^{4}He and ^{3}He(d,p)^{4}He fusion reactions[2] with first principle calculations based on realistic NN potentials. In addition, we delivered an evaluation of the (poorly known) n-^{3}H elastic cross section for 14 MeV neutrons,[3] important for plasma diagnostics in fusion experiments, with the required 5% accuracy.

More recently, we have taken initial steps at achieving a complete *ab initio* description of light-nucleus fusion and light exotic nuclei by incorporating the 3N force into our binary-reaction formalism[4] and extending our approach to the description of three-cluster dynamics.[5] These technical developments are highly nontrivial and individually pose an unprecedented computational challenge, let alone their combined treatment. Therefore, at first, we have been pursuing them separately. However, our objective for 2017 is to achieve the first *ab initio* description, complete with both 3N-force and three-cluster dynamic effects, of: 1) processes important for plasma diagnostics in fusion experiments and for the standard solar model, such as the ^{3}H(^{3}H,2n)^{4}He and ^{3}He(^{3}He,2p)^{4}He fusion reactions; and 2) the spectroscopy of exotic nuclei such as ^{11}Li, a ternary system of two nucleons orbiting around a ^{9}Li core, whose components are not bound in pairs. In addition, with increased access to HPC resources in 2017 we will be able to use our present theory and codes to perform high-fidelity calculations of the scattering of nucleons on a variety of p-shell targets, using NN+3N forces, by accessing much larger model spaces than possible today.

To achieve these goals in 2017 we will need a significantly larger amount of computational resources compared with our present usage. The concomitant storage of sparse matrices of large dimensions associated with the description of three-cluster dynamics and the enormous number (~ 6 billion) of 3N force matrix elements will dramatically increase the required amount of memory per node. Algorithmic changes will likely be required to further improve the distribution of matrix elements across all available processors and systematically implementing load balancing techniques that have proven very efficient in the most recently implemented parts of our production code.

3 Computational Strategies (now and in 2017)

3.1 Approach

Similar to *ab initio* nuclear structure calculations, which aim at describing the wave functions of nucleons bound inside a nucleus, we deal with a strongly correlated non-relativistic quantum many-body problem. We view the nucleons as point-like fermions, all of which interact among each other through high-precision two- and three-nucleon forces. However, in *ab initio* nuclear reactions we work

on the even harder problem of describing how nuclei – or clusters, as we like to call them – themselves interact with each other. To describe two-cluster (or three-cluster) dynamics, we use basis states made of pairs (or triplets) of nuclei in relative motion with respect to each other, in which each cluster of nucleons is described by its *ab initio* many-body wave function.

Mathematically our approach casts the many-body problem into one-dimensional (or two-dimensional, in the case of three-cluster dynamics) integral-differential coupled-channel equations for the cluster's relative motion. The main computational challenge is to construct the Hamiltonian and overlap (or norm) non-local couplings from the input two- and three-nucleon interactions and the many-body wave functions of the target and projectile nuclei. (The non-locality and the algebraic complexity of these couplings are a consequence of the Pauli exclusion principle, which is treated exactly in the formalism.) We use the R-matrix method on a Lagrange mesh to find the elements for the scattering matrix. At the end of the run we use the scattering matrix elements to calculate scattering and reaction observables to compare with available experimental data or to make needed predictions.

The Hamiltonian and the norm couplings can be obtained in second quantization in terms of matrix elements of creation and annihilation operators acting on the target nucleus. The number of nucleons affected (i.e., the rank of the operator involved) increases with the number of nucleons forming the projectile and going from two- to three-nucleon interactions. For example, while the description of a nucleon scattering on a nucleus with a two-nucleon interaction requires the evaluation of matrix elements of one and two-body densities of the target, calculations of deuterium-nucleus scattering including a three-nucleon force involve up to four-body densities. Storing in memory many-body density matrices is computationally very demanding. We adopt two computational strategies, which we have used to benchmark our results. One, more recent, strategy takes advantage of the target's implicit expansion in Slater determinants of single-nucleon states to perform an efficient on the fly calculation of the density matrix elements. In the second, first implemented, strategy we first calculate and store in memory the reduced matrix elements of tensor operators, from which we can then reconstruct the desired density on the fly. With the exception of the treatment of isospin symmetry, the two implementations are formally equivalent. The latter method is feasible only for very light systems but is more efficient for reactions with different projectiles, while the former is ideally suited for addressing heavier targets.

3.2 Codes and Algorithms

The calculations described above are performed using the NCSM_RGM (no-core shell model – resonating group method) reaction code, with inputs from the auxiliary TRDENS (transition densities) preprocessing code. Both are high-performance computing codes mostly dealing with sparse linear algebra, but also including components featuring dense linear algebra. NCSM_RGM uses an hybrid MPI/OpenMP parallelization to take full advantage of all computing cores as well as MPI/IO to mitigate I/O costs associated with large input files. In the newest routines, make use of dynamic load balancing techniques. The largest runs utilize 12,288 MPI tasks with 8 OpenMP threads per MPI task

(total size=98,304 cores, 6,144 new Jaguar XK6 nodes). TRDENS is based on the MPI-2 protocol and was run with up to 900 cores on SIERRA, a Dell supercomputing system with peak speed of 261 teraflop/s, twelve processors/node and 24GB memory per node.
Finally, for reactions involving heavier targets we are shifting large portions of the calculations to an auxiliary sister code, which we will call here NCSM_RGM_SD, that pre-computes the Hamiltonian and norm couplings within a Slater determinant single particle basis. NCSM_RGM_SD is also a hybrid code, i.e., adopting MPI parallelization between the individual compute nodes and OpenMP parallelization across the cores of each node. The largest runs have been performed on the NERSC Edison machine using 600 nodes, for a total of 14400 cores.

4 HPC Resources Used Today

4.1 Computational Hours

The LLNL Institutional Grand Challenge project "From Nucleons to Nuclei to Fusion Reactions", led by S. Quaglioni, has a significant allocation on SIERRA, a LLNL capability computer for moderate to large parallel jobs. SIERRA is a Dell supercomputing system, with peak speed of 261 teraflop/s, 12 processors/node and 24GB memory per node, for a total of 1,944 nodes. We also have access to CAB, an LLNL large capacity machine for small to moderate parallel jobs. CAB is a Linux cluster with 16 2.6GHz processors/node and 32GB memory per node for a total of 1,296 nodes. Our combined utilization in 2013 was over 10M Core-Hr.

About 23M Core-Hrs were also used in 2013 for reactions calculations on the Oak Ridge National Laboratory, Cray XK7 TITAN machine, as part of the INCITE project "Nuclear Structure and Nuclear Reactions", led by J. Vary.

Finally over the past year 2013 we made use of about 12M CPU hours on the Edison machine at NERSC to run the NCSM_RGM_SD code and calculate nucleon-nucleus scattering on p-shell nuclei, including the 3N force. This work was performed under the NERSC project "m94", led by J. Vary.

4.2 Parallelism

As of today the NCSM_RGM code has been run with as many as 98,304 cores (composed of 6,144 nodes, with 12,288 MPI tasks with 8 OpenMP threads per MPI task).

The TRDENS code was run with up to 900 cores on SIERRA.

The NCSM_RGM_SD code has been run on the Edison machine using between 200 and 600 nodes, for a total of 4800 to 14400 cores. Runs using twice the number of cores are likely feasible, given the good scaling performance presented by the code insofar.

Utilization for typical NCSM_RGM, TRDENS and NCSM_RGM_SD runs varies depending on the system under study (projectile/target mass, two- or three-nucleon Hamiltonian, size of the model space). Due to

limits on the size of runs per user, typical NCSM_RGM runs on LLNL machines are performed with 12,000 to 19,200 cores, and are restarted multiple times.

Our project requires strong scaling. The size of our problems is such that calculations would simply not be possible without parallel computing. We rely on scaling to complete our calculations, or major stages of it, within the allowed runtime.

4.3 Scratch Data

The required amount of temporary space varies somewhat with the problem under investigation. We require on the order of 1 TB per major run.

4.4 Shared Data

We have used the project folder within project 'm94' for the common use of certain matrix-element files for collaborations.

4.5 Archival Data Storage

We have not used the data archive.

5 HPC Requirements in 2017

5.1 Computational Hours Needed

To reach the scientific goals listed in section 2.2 as well as other scientific achievements that will be made possible by the sustained development of our *ab initio* reaction approach and codes, we forecast that in 2017 our project will require on the order of 75M Core-Hrs on a machine such as Edison.

We have not secured any significant allocation for 2017, although we expect that we will compete for high-performance computing allocations on LLNL and ORNL systems.

The primary factor driving the need for more core hours is the increased challenge of more complex reaction calculations including the 3N force on larger target and projectile nuclei.

5.2 Parallelism

In 2017 we expect that the NCSM_RGM code will typically use on the order of 28,400 to 60,000 conventional cores per run, with larger runs of up to 100,000 cores.

With further improvements on the MPI scaling, we expect that the NCSM_RGM code could be run with using the full Edison machine.

5.3 I/O

Our application has several built-in restart procedures, which allow us to perform fairly large calculations even on computers with a modest number of nodes such as SIERRA.

Our data requirement will grow in 2017, particularly in output. We expect that input should not exceed 100 GB, while we estimate that we will reach on the order of half a TB of output per run in our calculations of three-cluster dynamics. These output, including Hamiltonian and norm couplings, scattering phase shifts, scattering wave functions, as well as elements of the scattering matrix at various energy steps, can be moved to long term storage and utilized for later calculations of reaction observables.

We use MPI I/O to mitigate costs associated with large input files. We would like to keep the time devoted to I/O to less than 5% of the total runtime.

5.4 Future Data Needs (Please fill in the blanks)

In 2017, we expect to need 4 TB of temporary scratch disk space, 8 TB of NERSC project space (globally accessible shared data), and 30 TB of storage on NERSC HPSS. The growth in these requirements relative to 2013 is due primarily to the increased scale of the scientific problem under investigation. In particular we will have to store multiple instances of input three-nucleon force matrix elements, converged eigenvectors for the target, projectile and compound nuclei and transition densities, as well as converged scattering phase shifts, scattering wave functions, as well as elements of the scattering matrix at various energy steps.

5.5 Memory Required

We will need a minimum of 64 GB of memory per node, but would benefit from the maximum feasible memory per node , which would allow us to increase the model-space size and hence the fidelity of our calculations. We can use all available memory per node.

5.6 Emerging Technologies and Programming Models

For the time being we have not explored the use of GPUs, mainly because for our computational problem it is more straightforward to use OpenMP capabilities, and in part due to the lack of manpower and readily available knowhow to explore these new technologies. We envision that GPUs may be

advantageous during the solution of the two- and three-body dynamical equations, where the calculation of the scattering matrix requires dense linear algebra (matrix multiplications and inversions, eigenvalue problems) at each energy step.

Our software currently runs in production on Titan, but for the time being we do not take advantage of the GPUs. Rather, our NCSM_RGM software uses a hybrid MPI/OpenMP parallelization to take full advantage of all computing cores and memory in the nodes.

Currently we are not porting to, nor optimizing our codes for MIC architecture. Our strategy for exploiting GPUs and other new technologies will be to seek the help of experts in computational science, perhaps through the SciDAC NUCLEI collaboration, which has helped improve the ab initio nuclear structure code MFDn.

In the transition to new and emerging architectures the assistance of NERSC experts on the new architectures will be fundamental. Ideally, large allocations on NERSC machines should be accompanied by dedicated support from an expert on the computer architecture to help nuclear scientist port their codes and use the machine in the optimal way, maximizing the scientific impact of their calculations. This type of assistance should be accompanied by DOE programs such as SciDAC, that allow the close collaboration of nuclear physicists with computational scientists to solve compelling nuclear science problems. In addition, it would also be desirable to create positions at the intersection between computational science and nuclear physics that could foster the growth of a new generation of nuclear computational scientists.

5.7 Software Applications and Tools

We need to link our codes to the multithreaded Intel Math Kernel Library (MKL). Our compiler of preference is the Intel Fortran compiler. We make use of Lustre file system software to stripe large input files onto multiple hard-drives and speed up MPI I/O. At the moment we have implemented our own binary file formatting through MPI I/O, but we may transition to HDF5 in the next years.

5.8 HPC Services

We would benefit from consulting and account support, data analytics, training, and collaboration tools.

5.9 Additional Data Intensive Needs

We would need to be able to transfer data among collaborators in the USA, Canada and Europe. Currently we do not have a data management plan.

5.10 Additional Data Intensive Needs: Burst Buffer

Technology aimed at improving I/O would be useful for our work, although our data needs may be not as intensive as in other projects.

5.11 What Else?

6 Requirements Summary Worksheet

Please try to fill out this worksheet, based on your answers above, to be best of your ability prior to the review.

	Used at NERSC in 2013	Needed at NERSC in 2017
Computational Hours	~12M (Edison early users)	75M
Typical number of cores* used for production runs	4,800 to 14,400 (3% to 10% of machine)	28,800 to 60,000 (20% to 45% of machine)
Maximum number of cores* that can be used for production runs	28,800 (20% of machine)	Full machine
Data read and written per run	<1TB	< 1TB
Maximum I/O bandwidth	Not known	Not known
Percent of runtime for I/O	<5%	<5%
Scratch File System space	1 TB	4 TB
Shared filesystem space	2 TB	8 TB
Archival data	Not used	30 TB
Memory per node	All available memory	All available memory
Aggregate memory	<38 TB	Typical: 160 TB

* "Conventional cores." For GPUs and accelerators, please fill out section 4.7.

7 Additional Storage and I/O Questions

These questions are optional but your answers will provide additional useful data for NERSC. If you don't know the answer to any of these leave them blank.

For Scratch data (like Question 5.4):

- We have both serial and parallel I/O. It is parallel for the large files, but serial for smaller ones.

---

[1] P. Navrátil, S. Quaglioni and R. Roth, Phys. Lett. **B704**, 379 (2011).
[2] P. Navrátil and S. Quaglioni, Phys. Rev. Lett. **108**, 042503 (2012).
[3] J. A. Frenje *et al.,* Phys. Rev. Lett. **107**, 122502 (2011); P. Navrátil *et al*., LLNL-TR-423504 (2010); J. D. Anderson *et al.*, LLNL-TR-435981 (2010).
[4] G. Hupin, J. Langhammer, P. Navrátil, S. Quaglioni, A. Calci, And R. Roth, Phys. Rev. C **88**, 054622 (2013).
[5] S. Quaglioni, C. Romero-Redondo, P. Navrátil, Phys. Rev. C **88**, 034320 (2013); C. Romero-Redondo, S. Quaglioni, P. Navrátil, and G. Hupin, Phys. Rev. Lett. 113, 032503 (2014).