ANL/ALCF/ESP-17/10

Argonne NATIONAL LABORATORY

Quantum Monte Carlo Calculations in Nuclear Theory

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

ALCF Early Science Program (ESP) Technical Report

ESP Technical Reports describe the code development, porting, and optimization done in preparing an ESP project's application code(s) for the next generation ALCF computer system. This report is for a project in the Theta ESP, preparing for the ALCF Theta computer system.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

ANL/ALCF/ESP-17/10

Quantum Monte Carlo Calculations in Nuclear Theory

Technical Report for the ALCF Theta Early Science Program

edited by
Timothy J. Williams and Ramesh Balakrishnan

Argonne Leadership Computing Facility

prepared by
Steven C. Pieper, Alessandro Lovato, Ewing Lusk, Maria Piarulli, and Robert Wiringa

September 2017

This work was supported in part by the Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

Quantum Monte Carlo Calculations in Nuclear Theory

Steven C. Pieper*[1], Alessandro Lovato†[1], Ewing Lusk‡[2], Maria Piarulli§[1], and Robert Wiringa¶[1]

[1] Physics Division, Argonne National Laboratory, Argonne, IL
[2] Mathematics and Computing Science Division, Argonne National Laboratory, Argonne, IL

1 Introduction

This report is for a Tier-2 ESP which was granted 10,000,000 core hours for code tuning to enable future proposal submissions. As such the original science program could not be undertaken. We demonstrated good OpenMP and MPI scaling; this was important for our successful ALCC proposal and our pending INCITE proposal. We also did some new science work that has resulted in a paper submitted to Physical Review Letters.

2 Science Summary

To a remarkable extent, atomic nuclei can be described as collections of point-like particles whose dynamics are dictated by a nonrelativistic Hamiltonian involving two- and three-nucleon potentials. Ab initio approaches are aimed at solving the many-body Schrödinger equation associated with the nuclear Hamiltonian, which is a highly non-trivial problem because of the nonperturbative nature and the strong spin-isospin dependence of the nuclear forces. Within this framework it is possible to disentangle the theoretical uncertainty coming from modeling the nuclear potential and currents from that due to the approximations that are usually inherent in many-body techniques.

Our many-body method of choice is quantum Monte Carlo (QMC), in particular Green's function Monte Carlo (GFMC), which allows solving the nuclear Schrödinger equation with the required 1% accuracy level for both the ground- and the low-lying excited states of $A \leq 12$ (A is the number of nucleons) nuclei. Since 2000 we have been achieving excellent agreement with experiment for nuclei with increasing number of nucleons by using the Argonne V18 two-nucleon interaction and the Illinois three-nucleon interactions (AV18+IL7). During this period many other groups have been developing chiral effective field theory (χEFT) potentials; these are generally nonlocal and not suitable for QMC methods.

*spieper@anl.gov
†lovato@anl.gov
‡lusk@mcs.anl.gov
§mpiarulli@anl.gov
¶wiringa@anl.gov

We and collaborators have recently developed a local χEFT potential including the Δ-isobar degrees of freedom [4, 3] which can be used with QMC. The most immediate advantage of an explicit account of the Δ is the numerical consistency between the values of the low-energy constants inferred from either πN or NN scattering. Such a theory also naturally leads to the Fujita-Miyazawa three-nucleon forces as well as to two-nucleon electroweak currents. The inclusion of N^{3}LO contact terms allowed a fit to the Granada database of pp and np scattering up to 200 MeV laboratory energy with $\chi^2 \simeq 1.4$ for all the cutoff values considered. We used Theta to make GFMC calculations with one of these potentials and found agreement with data that is comparable to the AV18+IL7 results [2].

3 Codes, Methods and Algorithms

Our application code is GFMC, which has been in continuous production use for many years. It has steadily undergone development to take advantage of each new generation of parallel machine and was one of the first to deliver new scientific results each time. GFMC provides *ab initio* calculations of properties of light nuclei by solving the many-body Schrödinger equation using realistic two- and three-nucleon interactions. It is generally recognized as the most reliable method for nuclei with 6 to 12 nucleons.

GFMC begins with the construction of a trial wave function Ψ_T that is a symmetrized product of two- and three-body correlation operators acting on an antisymmetric A-body single-particle wave function that has the proper quantum numbers for the state of interest. The expectation value of the Hamiltonian H is evaluated by standard Metropolis Monte Carlo sampling in configuration space and the parameters in Ψ_T are varied to minimize the energy expectation value $\langle\Psi_T|H|\Psi_T\rangle$. GFMC then projects out the lowest eigenstate Ψ_0 of given quantum numbers starting from Ψ_T by a propagation in imaginary time τ: $|\Psi_0\rangle = \lim_{\tau\to\infty} \exp[-(H - E_0)\tau]|\Psi_T\rangle$. The propagation is carried out as a series of many small steps $\Delta\tau$. Expectation values of operators are evaluated as mixed matrix elements $O(\tau) = \langle\Psi_T|O|\Psi(\tau)\rangle$, and the behavior as a function of τ analyzed to obtain converged results. Because H and $\exp[-(H - E_0)\tau]$ commute, the mixed estimate is the exact expectation of $\langle\Psi(\tau/2)|O|\Psi(\tau/2)\rangle$ but linear extrapolations are used to evaluate other quantities.

The main part of GFMC is written in Fortran and compiles with modern Intel, gfortran, and IBM XL compilers. It makes extensive use of Fortran-90 features such as pointers, array subscripts, and defined types. It also uses older features such as common and equivalence to obtain efficiency. We use OpenMP to parallelize our Fortran code across cores within an MPI rank and to take advantage of hyperthreading within cores.

We make somewhat extensive use of the C preprocessor to enable one source file to generate subroutines, with different names, that are used in the same executable. This involves substituting different strings for a given string in the source. This is not supported in the Fortran standards but is available with both Intel and IBM compilers. As mentioned below, a work around was necessary for the Cray compiler.

GFMC relies on two libraries, ADLB and DMEM. ADLB [1] handles computational load balancing by accepting work packages from any MPI rank and distributing them to ranks that need work to do. ADLB also routes the work package answers back to the originating rank. DMEM carries out memory load balancing by storing large arrays on any node with enough memory and

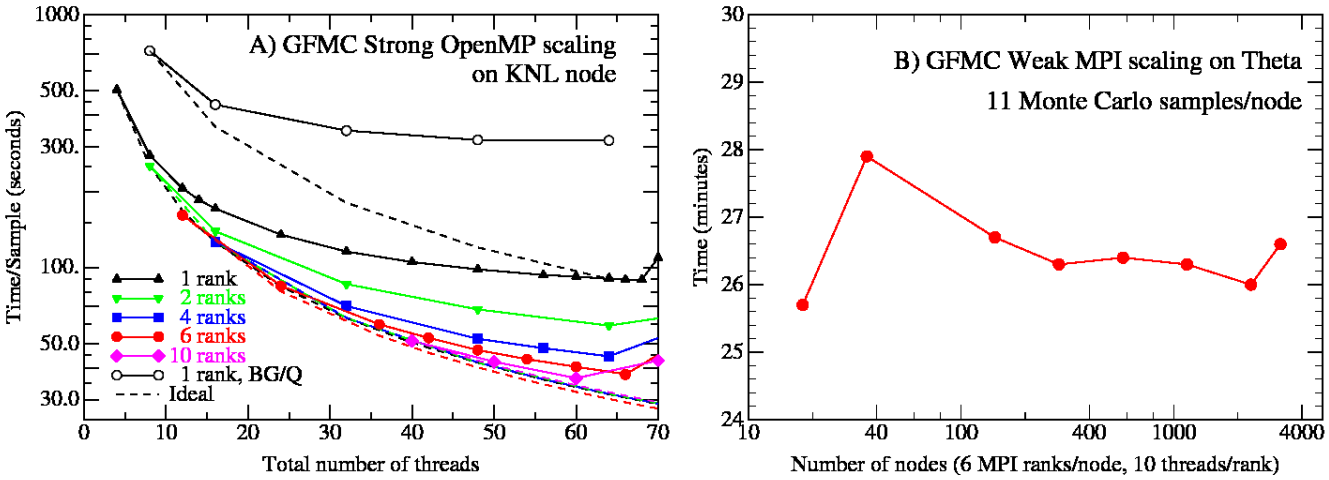

Figure 1: Scaling of GFMC calculations on KNL nodes and Theta

subsequently fetching them when needed. Both libraries are written in extremely portable C and developed locally by coinvestigator Lusk and collaborators, who continue to maintain them for the HPC community under SciDAC grants. ADLB and DMEM rely on MPI for interaddress space communication. Currently the only features of MPI-2 that we need are related to thread safety and I/O; the communication is carried out by using original MPI-1 functions.

4 Code Development

4.1 Optimization for Theta

As shown in Fig. 1, we have good MPI scaling to essentially the full Theta machine and good OpenMP scaling to the number of threads (cores) available to a given MPI rank. Panel A of the figure shows GFMC strong OpenMP scaling on a KNL node with 1 to 10 MPI ranks running on the node. We see good scaling for six to ten ranks up to using the full 64 cores on the node. Panel B shows weak MPI scaling on increasing numbers of nodes on Theta. Six ranks with 10 threads each were used on each node. The scaling is very good up to 3168 nodes, nearly the full machine (note the expanded time axis). Several OpenMP errors and inefficiencies using the Intel compiler with the Cray OS were found and reported by us during this work.

The greatest room for improvement in the speed of the the GFMC code is in the performance of a single core. Most of the GFMC time for large nuclei is spent in two types of operation. The first uses large tables of indices and coefficients to make the trial wave function vectors. These tables are generated before the calculation and are read from disk. The second operation is an iterative sequence of large sparse complex matrix multiplies on these vectors. The matrices consist of noncontiguous 4×4 and 8×8 blocks. The multiplies are coded as specific table-driven Fortran subroutines that make use of the detailed structure of these blocks [5]. We found that the AVX512 instructions do not substantially increase the speed of these operations. The compiler optimization reports suggest that this is because we are working with COMPLEX(8) arrays which, by the Fortran standard, have pairs of real and imaginary numbers. Thus they do not meet the AVX requirement

of having unit stride. We get no significant AVX improvement whether we declare the arrays to be COMPLEX(8) or REAL(8) with an extra leading dimension of 2.

Recently we have been using Argonne's new Bebop computer which has both Broadwell and Knights Landing nodes. We find that we get close to twice the throughput on the Broadwell nodes than on the KNL nodes.

4.2 Other Changes

The program was already making good use of OpenMP and (through ADLB and DMEM) MPI on both Mira and Argonne's Blues computers. We did make some further improvements in DMEM (which had only recently been developed). We added subroutines to compute the new χEFT potentials.

Our use of the C preprocessor for Fortran source is not supported by the Cray compiler. Instead of having the Makefile directly invoke the compiler, we had to have it invoke a script that ran the Fortran through the C compiler and captured the modified Fortran source. This is then given to the Cray Fortran compiler.

5 Portability

The program was already working on Mira (IBM BG/Q) and Argonne's Blues (Intel Sandybridge). There were no significant problems getting it to work (asside from some efficiency issues) on Theta using the Intel compiler. However significant changes had to be made to the Makefiles to use the Cray compiler because of our use of the C preprocessor.

6 Conclusions

Although the limited CPU hours provided by this Tier-2 grant did not allow us to address the original science question, the grant was very valuable in allowing us to learn how to use Theta and prepare subsequent proposals. We also did run calculations for a more limited science objective and a paper based on those results is now submitted.

References

[1] E. Lusk, S. C. Pieper, and R. Butler. More scalability, less pain: A simple programming model and its implementation for extreme computing. *SciDAC Review*, 17:30, 2010.

[2] M. Piarulli, A. Baroni, L. Girlanda, A. Kievsky, A. Lovato, Ewing Lusk, L. E. Marcucci, Steven C. Pieper, M. Schiavilla, R. andViviani, and R. B. Wiringa. Local chiral potentials with Δ intermediate states and the structure of light nuclei. *arXiv*, 1707.02883, 2017.

[3] M. Piarulli, L. Girlanda, R. Schiavilla, A. Kievsky, A. Lovato, L. E. Marcucci, Steven C. Pieper, M. Viviani, and R. B. Wiringa. Local chiral potentials with Δ intermediate states and the structure of light nuclei. *Phys. Rev. C*, 94(5):054007, 2016.

[4] M. Piarulli, L. Girlanda, R. Schiavilla, R. Navarro Prez, J. E. Amaro, and E. Ruiz Arriola. Minimally nonlocal nucleon-nucleon potentials with chiral two-pion exchange including Δ resonances. *Phys. Rev. C*, 91(2):024003, 2015.

[5] S. C. Pieper. Quantum Monte Carlo Calculations of Light Nuclei. In A. Covello, F. Iachello, R. A. Ricci, and G. Maino, editors, *Proceedings of the "Enrico Fermi" Summer School, Course CLXIX, Nuclear Structure far from Stability: New Physics and new Technology*, page 111. IOS Press, Amsterdam, 2008. Reprinted in La Rivista del Nuovo Cimento, **31**, 709, (2008).

Argonne Leadership Computing Facility
Argonne National Laboratory
9700 South Cass Avenue, Bldg. #240
Argonne, IL 60439

www.anl.gov

Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC