Computing Atomic Nuclei on the Cray XT5

Hai Ah Nam and David J. Dean
Oak Ridge National Laboratory
James P. Vary and Pieter Maris
Iowa State University

Abstract

Understanding the structure of atomic nuclei is crucial for answering many fundamental questions such as the origin of elements in the universe. The computational challenge for large-scale nuclear structure calculations, particularly *ab initio* no-core shell model calculations (NCSM), stems from the memory constraints of the Hamiltonian matrix, with dimensions in the billions. We push the limits of previous light nuclei NCSM nuclear structure calculations to study carbon-14, commonly used in radioactive carbon dating. The principal goal of this study is to address important questions about the low-lying structure of ^{14}C that leads to it anomalously, but very useful, long lifetime. Here, we discuss the issues of scaling MFDn, a nuclear shell model application, to study ^{14}C and the new I/O demands on the CrayXT5.

KEYWORDS: nuclear shell model, carbon-14, XT5

1 Introduction

Radioactive dating, specifically of organic materials, relies on the anomalously long half-life of ^{14}C, which is approximately 5,730 years. ^{14}C, with 6 protons and 8 neutrons, decays via beta emissions, whereby a neutron in the nucleus becomes a proton, an electron, and an antineutrino

$$\text{beta-decay: } n \longrightarrow p + e^- + \overline{\nu}_e, \quad (1)$$

$$\text{thus, } {}^{14}C \longrightarrow {}^{14}N + e^- + \overline{\nu}_e \quad (2)$$

resulting in ^{14}N, with 7 protons and 7 neutrons. The long lifetime of this isotope is primarily due to detailed nuclear structure effects in the low-lying energy states that lead to an anomalously small Gamow-Teller transition probability (B(GT)) given by

$$B(GT) = \frac{1}{2J_i + 1}|M_{GT}|^2, \quad (3)$$

where J_i is the total angular momentum of the initial state and M_{GT} is the Gamow-Teller matrix element calculated from the wavefunctions describing the nucleus. The half-life ($\tau_{1/2}$) is inversely proportional to the B(GT) value, thus a long half-life requires a small Gamow-Teller matrix element.

Studies, starting as early as the 1950's, have yet to result in a satisfactory description for the reason for the long half-life. From a nuclear structure point of view, this is a challenging problem that has resisted a purely microscopic understanding, and has led to the invocation of in-medium modifications of the nuclear force. Recent advances in computing technology and nuclear theory allow us to finally study ^{14}C from a fundamental (*ab initio*) formulation. Using Jaguar XT5, we can perform a comprehensive *ab initio* study in the No-Core Shell Model (NCSM) framework, utilizing realistic interactions that properly include spin-orbit effects induced by three-nucleon interactions.

Although Jaguar XT5 provides unmatched computational resources, the memory requirements for these calculations exceeds the available memory of the system. To attain the highest precision calculation currently viable, additional code modifications are necessary to contend with the memory demands of an *ab initio* approach. Using the Many Fermion Dynamics - nuclear (MFDn) code, we describe code improvements in progress involving out-of-core techniques. These improvements will allow for the first large-basis calculations of ^{14}C and ^{14}N, with basis dimensions in the billions, using an *ab initio* three-nucleon interaction in the NCSM. We also present here some preliminary results from smaller model space calculations to motivate the necessity for higher model space calculations.

2 The Mystery of ^{14}C

Allowed decays occur with higher probability when conditions on the correct quantum numbers describing the transition are met. The transition from the ground state of ^{14}N to the ground state of ^{14}N meet the conditions for beta-decay transitions, thus should result in a very short half-life. But, we know that, in fact, the half-life is quite long. This means the Gamow-Teller transition strength (B(GT)) should be small and the GT matrix element very close to zero ($\sim 10^{03}$), not seen in the current theoretical calculations. Also puzzling in the structure of ^{14}C is the evidence of charge-symmetry breaking in that the long half-life of ^{14}C (6 protons, 8 neutrons) cannot be explained by decay of its analog nucleus ^{14}O (8 protons, 6 neutrons). Figure 1 shows that ^{14}O has a short half-life, on the order of seconds, along with other light nuclei that undergo beta decay. Strangely, only two light nuclei, ^{10}Be and ^{14}C, have half-lives on the order of years.

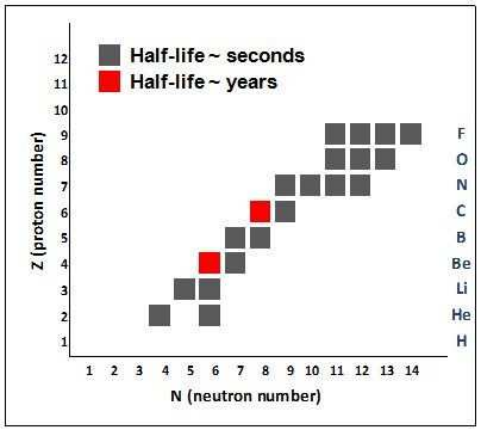

Figure 1: Chart of light nuclei that decay via beta emissions.

Calculations for the B(GT) values for the transition of ^{10}B to ^{10}Be [6] using 2- and 3-body interactions with three-nucleon forces in the NCSM framework show good agreement with experiment. Similar calculations for the ^{14}N to ^{14}C transition using 2-body interactions [1, 7] have yet to provide satisfactory agreement, requiring additional modifications to the nuclear interaction to achieve good results [2]. A systematic and comprehensive *ab initio* study of ^{14}C utilizing 2- and 3-nucleon forces is in order to properly determine the spin-orbit effects and its contribution to ^{14}C's anomalously long half-life.

2.1 Calculations

This study requires a series of large-basis shell model calculations for A=14 to examine the behavior of the GT transition matrix element on components of the nuclear Hamiltonian. These calculations are carried out within the framework of the *ab initio* no-core shell model (NCSM). At present, the NCSM is the only method capable for treating the A=14 system in an *ab initio* framework. We will provide a systematic study with model spaces extending to N_{max}=10 for two-body interactions and N_{max}=8 for three-body interactions. At N_{max}=10, the basis dimension approaches 20 billion. These will be the largest, and most challenging NCSM calculations performed to date.

First, we will study the behavior of the GT matrix elements with realistic two-nucleon interactions up to N_{max}=10, and then with the non-local two-body interactions JISP16 and INOY. These interactions are based on two-nucleon scattering data, but have realistic modifications to off-shell components that are designed to account for explicit three-nucleon (and possibly higher-body) components in the nuclear interaction. The advantage of these non-local two-nucleon potentials is that two-body interactions are computationally easier to handle, and offer a pathway to larger model spaces and perhaps better convergence. In addition, we also plan to carry out calculations in model spaces up to N_{max}=8 with explicit three-nucleon interactions based on effective-field theory. Even in this smaller space, these N_{max}=8 calculations are likely at least an order-of-magnitude more difficult than the two-body N_{max}=10 calculations. These calculations will offer a comprehensive study to determine sensitivity of the GT matrix elements on the model space and the various components of the nuclear Hamiltonian. By using four different interactions, we can also investigate the viability of using non-local two-body interactions to account for explicit three-body effects. No calculations including the effects of three-nucleon interactions have been performed yet.

3 Numerical Method

Solving the quantum many-body system is a common problem found throughout many areas of physics and chemistry. The No-Core Shell Model

(NCSM) is an *ab initio* method of solving the quantum many-body problem, with success in describing light nuclei ($A \leq 16$) [6, 4, 5, 3]. The NCSM solves the A-body Schrödinger equation

$$H\Psi(\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_k) = \lambda\Psi(\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_k) \qquad (4)$$

in a harmonic oscillator basis and diagonalizes the Hamiltonian matrix to find the eigenvalues and eigenvectors characterizing the energy states of the nucleus. Here H is the many-body Hamiltonian, Ψ is the many-body wavefunction with corresponding eigenvalue λ.

The shell model Hamiltonian using only the two-nucleon interaction is of the form

$$H = \frac{1}{A}\sum_{i<j}^{A} \frac{(\overrightarrow{p}_i - \overrightarrow{p}_j)^2}{2m} + \sum_{i<j}^{A} V_{NN,ij} \qquad (5)$$

where A is the total number of nucleons, m is the nucleon mass, $V_{NN,ij}$ is the two-nucleon interaction between particle i and j. The natural extension for inclusion of the three-nucleon interaction is

$$H = \frac{1}{A}\sum_{i<j}^{A} \frac{(\overrightarrow{p}_i - \overrightarrow{p}_j)^2}{2m} + \sum_{i<j}^{A} V_{NN,ij} + \sum_{i<j<k}^{A} V_{NNN,ijk} \qquad (6)$$

where $V_{NNN,ijk}$ is the three-nucleon interaction. A detailed description of the NCSM can be found in reference [6].

Diagonalization of the Hamiltonian matrix using the Lanczos algorithm gives us the low-lying energy spectrum (eigenvalues) and nuclear wavefunction (eigenvectors), which are necessary for computing observables and transition strengths, like the B(GT) values. The Lanczos algorithm [9] is an iterative eigenvalue solving method derived from the Arnoldi method, specifically for real symmetric matrices. It reduces the full matrix to tridiagonal form, making it easier to solve for the eigenvalues. This method is ideal for solving a large sparse matrix where only the extremum (smallest) eigenvalues are of interest. Current shell model codes require hundreds of iterations to achieve convergence for the lowest 10 eigenvalues. The primary operations include matrix-vector multiplications and vector dot products.

The difficulty of implementing the NCSM method is the Hamiltonian matrix which is large, sparse, and irregular for large-scale calculations of physical interest. Calculating properties of larger nuclei with more reliable results requires a large basis dimension, hence the size of the matrix, and including higher body forces, increasing the number of non-zero matrix elements. Scaling a shell model code for larger nuclei and including higher body forces poses a complex computational challenge due to the intense memory constraints.

3.1 Many Fermion Dynamics - nuclear (MFDn)

The MFDn code has been steadily developed over the last two decades [10]. It is a nuclear shell model application that stores the non-zero matrix elements onto memory for quick retrieval for the matrix-vector multiplies during each Lanczos iteration. Given that there are roughly 500 iterations to converge the lowest 10-15 eigenvalues, speed is a necessity for each iteration. Other options for contending with the large number of matrix elements include on-the-fly recalculation for each iteration as well as storage of the matrix elements onto disk, both of which are constrained by run-time as well as the memory for the data structures to recompute or retrieve the information.

Recent improvements through SciDAC collaborations [8] have greatly improved the performance of MFDn through efficient memory balancing of the matrix elements across the cores. This was achieved through distribution of the many-body basis states with a corresponding distribution for the Hamiltonian matrix elements seen in figure 2 and a cyclic distribution of the non-zero matrix elements. Due to these improvements, MFDn has scaled to 30,000 cores on Jaguar XT4. The biggest limitation to further scaling MFDn is the availability of cores, or more specifically, the availability of memory available to each core.

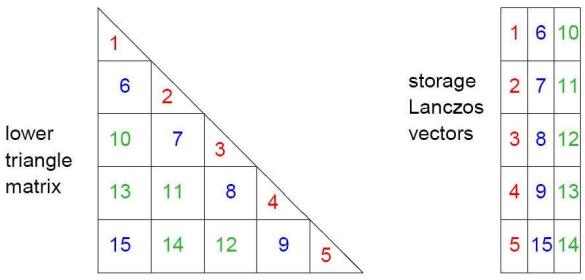

Figure 2: Distribution scheme of the lower triangular matrix and lanczos vector. The numbers represent processor IDs.

^{14}C					
	Basis	$\neq 0$ Matrix Elements		# of cores	
N_{max}	Dimension	2-body	3-body	2-body	3-body
2	5.80E+03	4.00E+05	2.90E+06	6	6
4	7.32E+05	1.62E+08	2.80E+09	28	190
6	3.37E+07	1.55E+10	4.42E+11	190	8128
8	8.73E+08	6.97E+11	2.90E+13	18336	148785
10	1.54E+10	1.94E+13	—	148785	—

^{14}N					
	Basis	$\neq 0$ Matrix Elements		# of cores	
N_{max}	Dimension	2-body	3-body	2-body	3-body
2	8.40E+03	7.00E+05	5.20E+06	6	6
4	9.75E+05	2.29E+08	4.10E+09	28	190
6	4.32E+07	2.07E+10	6.08E+11	496	10296
8	1.09E+09	9.01E+11	3.90E+13	18336	148785
10	1.89E+10	2.45E+13	—	148785	—

Figure 3: Dimensionality for ^{14}C and ^{14}N 2-body and 3-body calculations, including an estimated number of non-zero matrix elements and recommended number of cores (for less than 4 hour run-time).

4 Petascale Scaling Issues

Jaguar XT5, with roughly 150,000 cores, 1.382 peak petaflops, and 300 terabytes of memory, is an extraordinary computational resource. But, pitted up against the tremendous demands of the NCSM calculation, it falls just below the needs to complete the largest model space calculations of the Gamow-Teller transition strength for ^{14}C. Figure 3 shows the dimension of the basis, the estimated number of non-zero matrix elements and the recommended number of cores for a less than 4 hour run-time for 2- and 3-body calculations of ^{14}C and ^{14}N. To compute the B(GT) values, wavefunctions from both ^{14}C and ^{14}N are needed.

For example, to complete the largest 3-body calculation at $N_{max} = 8$ for ^{14}C and ^{14}N requires storage of roughly 3×10^{13} and 4×10^{13} matrix elements respectively. The memory required by each core to store its share of the matrix elements and indexing information, along with the input and output vectors is roughly given by $2\times4\times$(# of matrix elements)/(# of cores) + $5 \times 4\times$(Basis Dimension)/(diagonals) + .15GB overhead, where the number of diagonals is given by (# of cores) = diagonals(diagonals + 1)/2. This results in 1.71 GB and 2.25 GB of memory needed per core to complete the calculation with full storage of the matrix elements. Given that the usable memory per core is roughly 1.7 GB, completing this calculation requires modifications to the management of the matrix elements.

In order to fit the largest basis calculations we are currently implementing integer compression techniques to reduce the overall memory footprint as well as storing some of the matrix elements out-of-core. Out-of-core techniques are particularly challenging at the scale of Jaguar XT5, ∼149,000 cores. Given that each core needs to read hundreds of megabytes of data per lanczos iteration, when the file system provides a maximum throughput of 200 GB/sec, implies that run-time will also become a huge constraint. Completing these large-scale calculations will require a fine balance between memory and run-time, utilizing everything in our bag of tricks to squeeze ^{14}C into an XT5... really it's ^{14}N.

5 Results

The full calculation of the B(GT) requires wavefunctions from both ^{14}C and ^{14}N. As a preliminary gauge of the results, Figure 4 show the B(GT) transitions calculated within ^{14}N. The results from the three-nucleon interaction show the B(GT) values to be much more stable across various harmonic oscillator frequency values, a desired effect in shell model calculations. Also, we see that the B(GT) value is smaller using the three-nucleon interaction versus the two-nucleon interaction. These results are encouraging and indicate that higher model space calculations are necessary to for improved convergence.

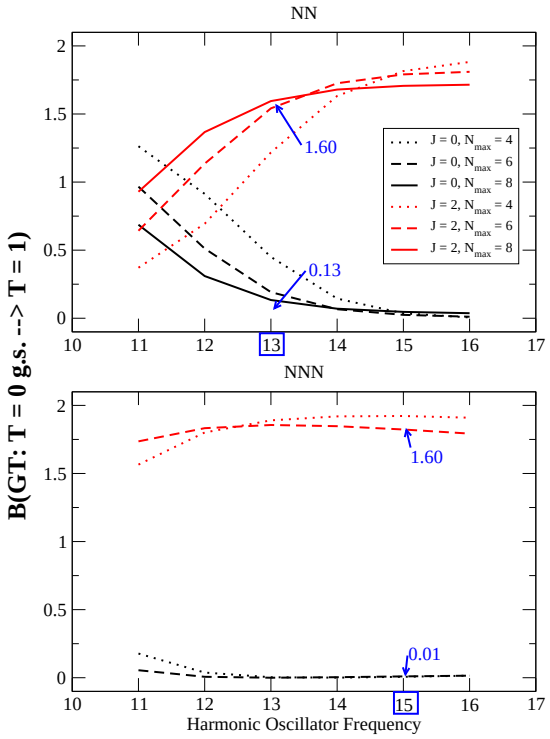

Figure 4: Preliminary results for B(GT) transitions within ^{14}N from the $(J^\pi, T) = (1^+, 0)$ ground state to the lowest $(0^+, 1)$ (black) and $(2^+, 1)$ (red) excited states of ^{14}N for specified model space N_{max}. The top panel uses a two-nucleon interaction and the bottom panel shows the three-nucleon calculations.

6 Conclusions

In conclusion, our preliminary results of lower model space calculations, up to $N_{max} = 8$ for 2-body and $N_{max} = 6$ for 3-body, show improvements using the three-nucleon interaction to the B(GT) transition values within ^{14}N. These results are encouraging for continuing development of the code to complete the largest model space calculations. Completion of this study with the highest model space calculations as well as final B(GT) values are necessary to determine the total effects of adding a three-nucleon force. This research is ongoing through the next couple of months under the Petascale Early Science period, made available at the National Center for Computational Sciences at ORNL on Jaguar XT5.

7 Acknowledgements

This research used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

8 About the Authors

Hai Ah Nam is a computational scientist at the National Center for Computational Sciences at Oak Ridge National Laboratory. Her research interests include low-energy nuclear structure research at the petascale and beyond. She can be reached at *namha@ornl.gov*.

David J. Dean is the Nuclear Theory Group Leader, Physics Division at Oak Ridge National Laboratory. He can be reached at *deandj@ornl.gov*.

James P. Vary is a professor in the Nuclear Physics Theory Group at Iowa State University. He is the author of the MFDn code. He can be reached at *jvary@iastate.edu*.

Pieter Maris is a postdoctoral researcher in the Nuclear Physics Theory Group at Iowa State University. He currently develops and maintains MFDn. He can be reached at *pmaris@iastate.edu*.

References

[1] S. Aroua, P. Navratil, L. Zamick, M.S. Fayache, B.R. Barrett, J. P. Vary, N. Mirnova, and K. Heyde, *Allowed gamow-teller excitations*

from the ground state of ^{14}n, Nuc. Phys. A **720** (2003), 71–83.

[2] J.W. Holt, G.E. Brown, T.T.S. Kuo, J.D. Holt, and R. Machleidt, *Shell model description of the* ^{14}c *dating* β *decay with the brown-rho-scaled NN interactions*, Phys. Rev. Lett. **100** (2008), 062501.

[3] P. Maris, J. P. Vary, and A. M. Shirokov, *Ab initio no-core full configuration calculations of light nuclei*, Phys. Rev. C **79** (2009), 014308.

[4] P. Navratil, V. G. Gueorguiev, J. P. Vary, W. E. Ormand, and A. Nogga, *Structure of* $a = 10-13$ *nuclei with two- plus three-nucleon interactions from chiral effective field theory*, Phys. Rev. Lett. **99** (2007), 042501.

[5] P. Navratil, V. G. Gueorguiev, J. P. Vary, W. E. Ormand, A. Nogga, and S. Quaglioni, *Light nuclei from chiral eft interactions*, Few-Body Systems **0** (2008), 1–6.

[6] P. Navratil and W. E. Ormand, *Ab initio shell model with a genuine three-nucleon force for the p-shell nuclei*, Phys. Rev. C **68** (2003), 034305.

[7] A. Negret, T. Adachi, B.R. Barrett, C. Bumer, A.M. van den Berg, G. P. A. Berg, P. von Brentano, D. Frekers, D. De Frenne, H. Fujita, K. Fujita, Y. Fujita, E.-W. Grewe, P. Haefner, M. N. Harakeh, K. Hatanaka, K. Heyde, M. Hunyadi, E. Jacobs, Y. Kalmykov, A. Korff, K. Nakanishi, P. Navrtil, P. von Neumann-Cosel, L. Popescu, S. Rakers, A. Richter, N. Ryezayeva, Y. Sakemi, A. Shevchenko, Y. Shimbara, Y. Shimizu, Y. Tameshige, A. Tamii, M. Uchida, J. Vary, H. J. Wrtche, M. Yosoi, and L. Zamick, *Gamow-teller strengths in the a=14 multiplet: A challenge to the shell model*, Phys. Rev. Lett. **97** (2006), 062502.

[8] P. Sternberg, E.G. Ng, C. Yang, P. Maris, J.P. Vary, M. Sosonkina, and H.V. Le, *Accelerating configuration interaction calculations for nuclear structure*, Proceedings of Supercomputing 2008 (2009).

[9] L. N. Trefethen and D. Bau, *Numerical linear algebra*, SIAM, Philadelphia, 1997.

[10] J.P. Vary, *The many-fermion dynamics shell-model code*, Iowa State University (1992), unpublished.