Physics capabilities and results with Many Fermion Dynamics – nuclear physics

Pieter Maris
pmaris@iastate.edu

Collaborators

Nuclear physics

ISU: James Vary, PM, students
LLNL: Petr Navratil, Erich Ormand
SDSU: Calvin Johnson, Plamen Krastev
ORNL/UT: David Dean, Hai Ah Nam,
Gaute Hagen, Thomas Papenbrock
OSU: Dick Furnstahl, students
MSU: Scott Bogner
WMU: Mihai Horoi
LANL: Joe Carlson
UA: Sid Coon, Bira van Kolck,
Bruce Barrett, Alex Lisetskiy, students
LSU: Jerry Draayer, students

International Collaborators:

Shirokov, Forssen, Schwenk,
Abe, Otsuka, Shimizu, Utsuno

Computer Science/Applied Math

Ames Lab: Masha Sosonkina, students
LBNL: Esmond Ng, Chao Yang, Philip Sternberg
ANL: Stefan Wild
OSU: Umit Catalyurek

Many Fermion Dynamics (MFD) – nuclear physics

Numerical approach for bound states of nucleons
using basis-space expansion methods

- Given
  - a V_{NN} and V_{NNN} (and V_{NNNN}) interaction
  - number of protons and neutrons: Z, N

  calculates
  - bound state spectrum and corresponding wave functions
  - one-body density matrices (in development)
  - selected observables:
    rms radii, magnetic moments, quadrupole moments,
    transition rates between states within the same nucleus
- Other applications
  - Hadron physics:
    relativistic field theory, number of particles not conserved
  - Atomic physics: trapped bosons and/or fermions

MFD_nuclear

- Platform independent MPI Fortran code
- Variational: for any finite truncation of the basis space, MFDn gives an upper bound for the ground state energy
- Smooth approach to asymptotic values with increasing basis space: No-Core Full Configuration calculation
- Convergence: Physics results independent of truncation ($N_{\max}$) and basis space parameter ($\hbar\omega$)
  - different methods (NCFC, CC, GFMC, . . .) using the same interaction give same results within numerical errors
- Single most important computational issue: exponential increase of the dimensionality with increasing H.O. levels

Dimensions and sparsity of matrices

Estimates of aggregate memory needed for storage of sparse symmetric Hamiltonian matrix in compressed column format

nucleus	$N_{\max}$	dimension	2-body	3-body	4-body
^{6}Li	12	$4.9 \cdot 10^{6}$	0.6 GB	33 TB	590 TB
^{12}C	8	$6.0 \cdot 10^{8}$	4 TB	180 TB	4 PB
^{12}C	10	$7.8 \cdot 10^{9}$	80 TB	5 PB	140 PB
^{16}O	8	$9.9 \cdot 10^{8}$	5 TB	300 TB	5 PB
^{16}O	10	$2.4 \cdot 10^{10}$	230 TB	12 PB	350 PB
^{8}He	12	$4.3 \cdot 10^{8}$	7 TB	300 TB	7 PB
^{11}Li	10	$9.3 \cdot 10^{8}$	11 TB	390 TB	10 PB
^{14}Be	8	$2.8 \cdot 10^{9}$	32 TB	1100 TB	28 PB
^{20}C	8	$2 \cdot 10^{11}$	2 PB	150 PB	6 EB
^{28}O	8	$1 \cdot 10^{11}$	1 PB	56 PB	2 EB

(presented at *Extreme Scale Computing Workshop – nuclear physics* Washington DC Jan 2009)

MFD_nuclear – recent developments

- Efficient construction of large sparse many-body matrix

  Sternberg, Ng, Yang, Maris, Vary, Sosonkina, Le, *Accelerating Configuration Interaction calculations for nuclear structure*, presented at SuperComputing 2008

- Parallel IO (MPI_IO, pHDF5)

  Laghave, Sosonkina, Maris, Vary, *Benefits of Parallel IO in Ab Initio Nuclear Physics Calculations*, presented at ICCS 2009, Baton Rouge, LA

- External field (2008) new results for neutron droplets in progress

- Generalized truncation schemes (2008/2009)

  collaboration with Abe, Otsuka, Shimizu, Utsuno with support from JUSTIPEN

- Generalized basis functions (2008/2009) Negoita, PhD thesis

- Scattering applications (2008/2009) Shirokov *et.al.*

- Total-J basis, in addition to usual m-scheme basis (2009)

  serial code functional, parallel version under development

- Hybrid MPI/OpenMP code (2009)

- Partial “regenerate-on-the-fly” (2009)

  under development, main purpose: reduce memory requirements

IO Challenges

- Input interaction files
  - need to be read by all processors
  - 3-body, $N_{\max} = 8$ for p-shell nuclei: 33 GB
  - current solution (2009): MPI_IO

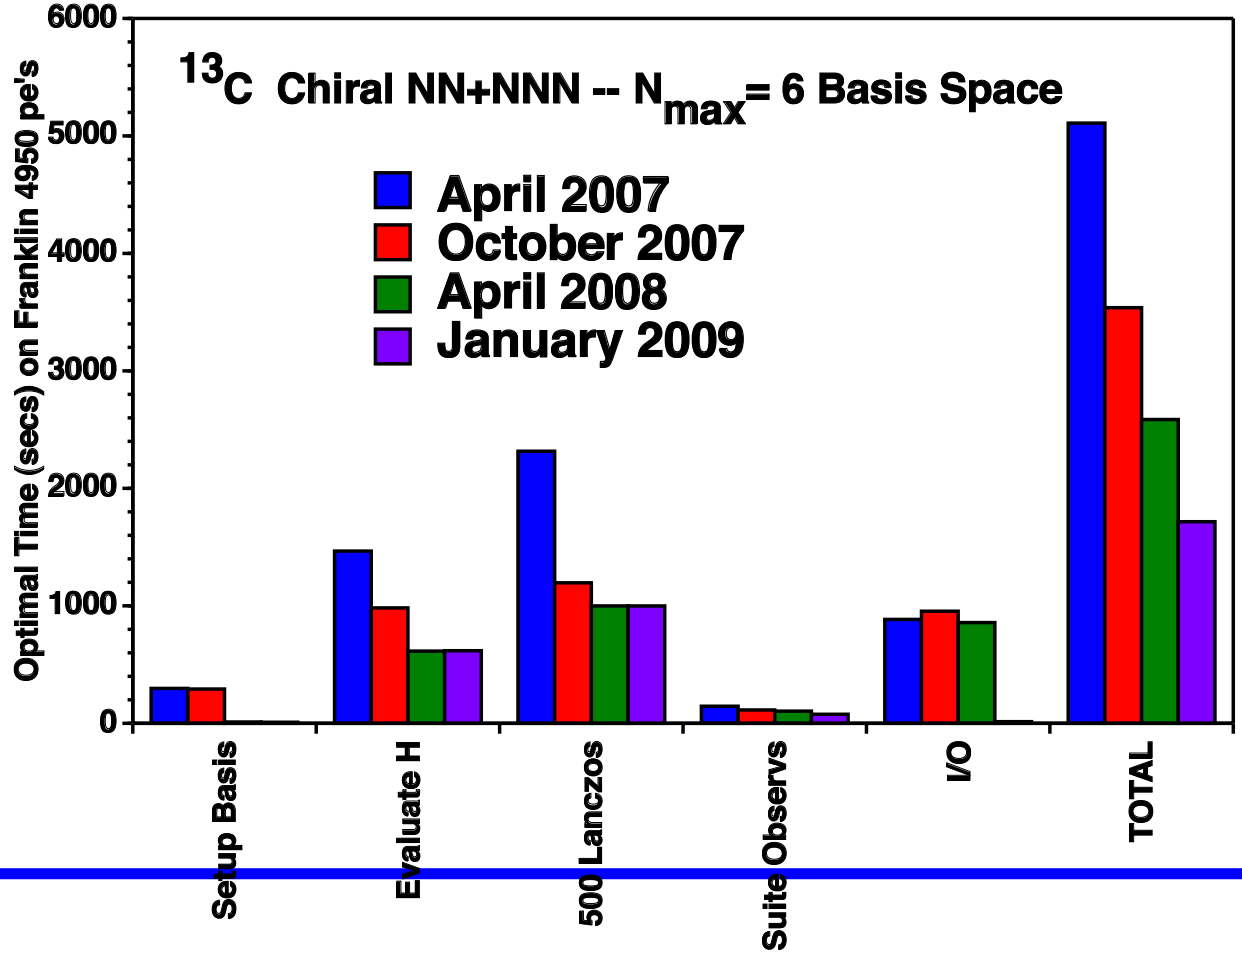

dimension $38 \cdot 10^6$
nonzero m.e. $56 \cdot 10^{10}$
size input 3 GB

IO Challenges

- Input interaction files
- Output
  - selected physics observables: ASCII by root
  - wavefunctions: parallel IO using “diagonal” processors

lower triangle matrix

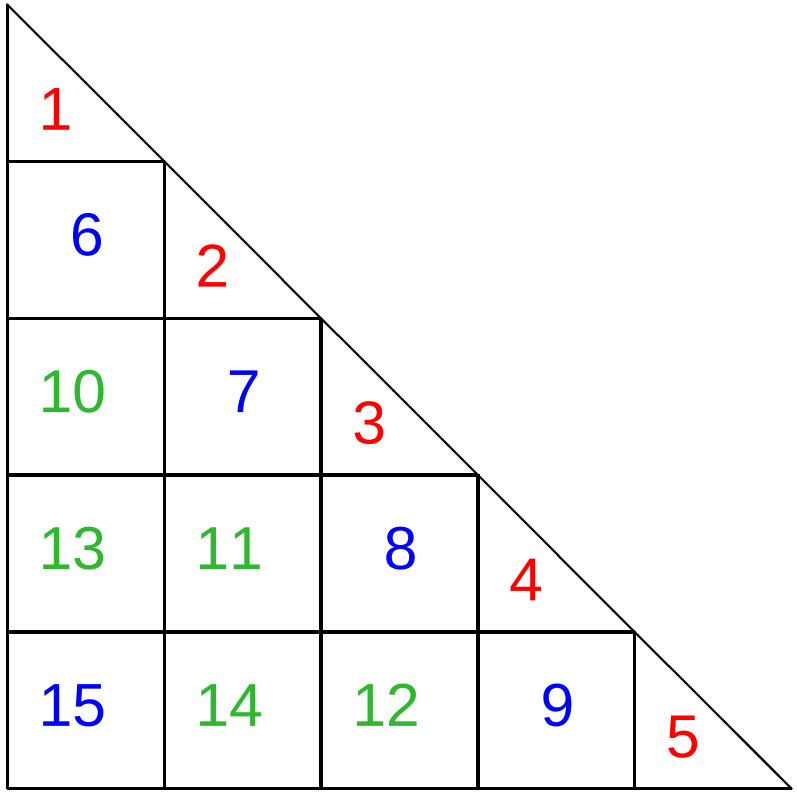

storage Lanczos vectors

1	6	10
2	7	11
3	8	12
4	9	13
5	15	14

IO of wavefunctions – presented at ICCS 2009

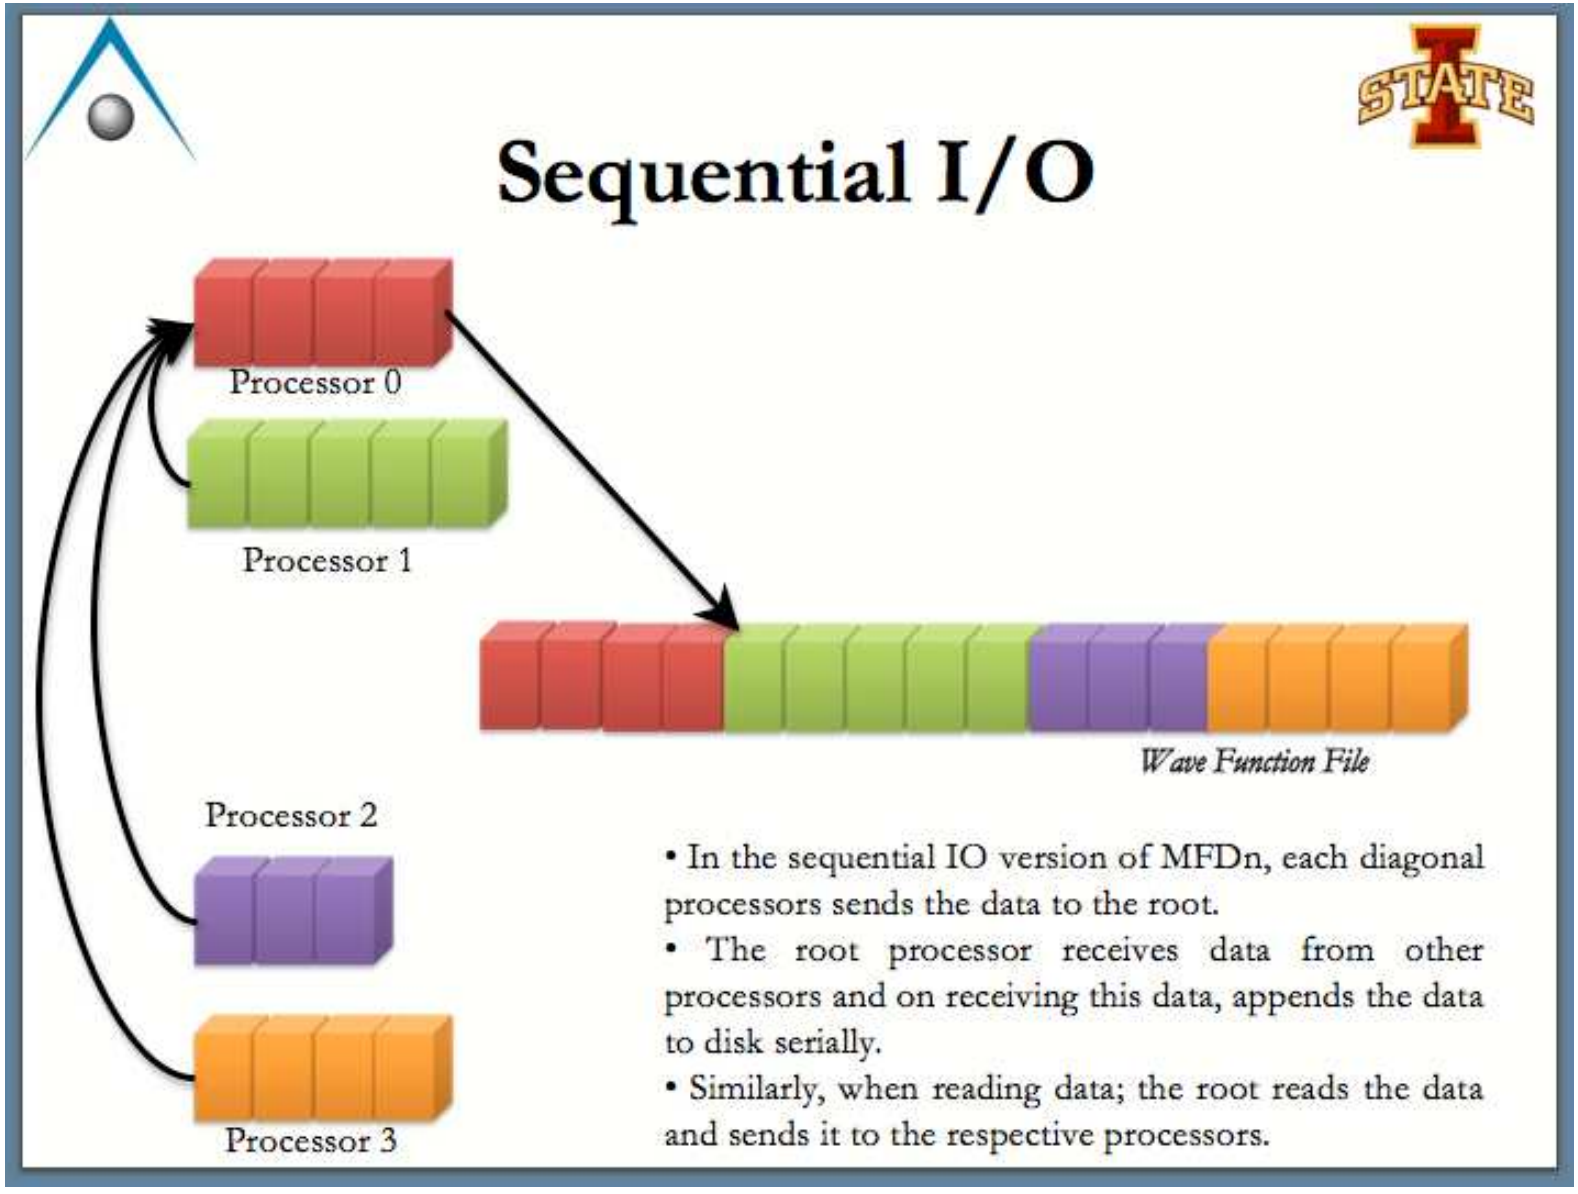

IO of wavefunctions – presented at ICCS 2009

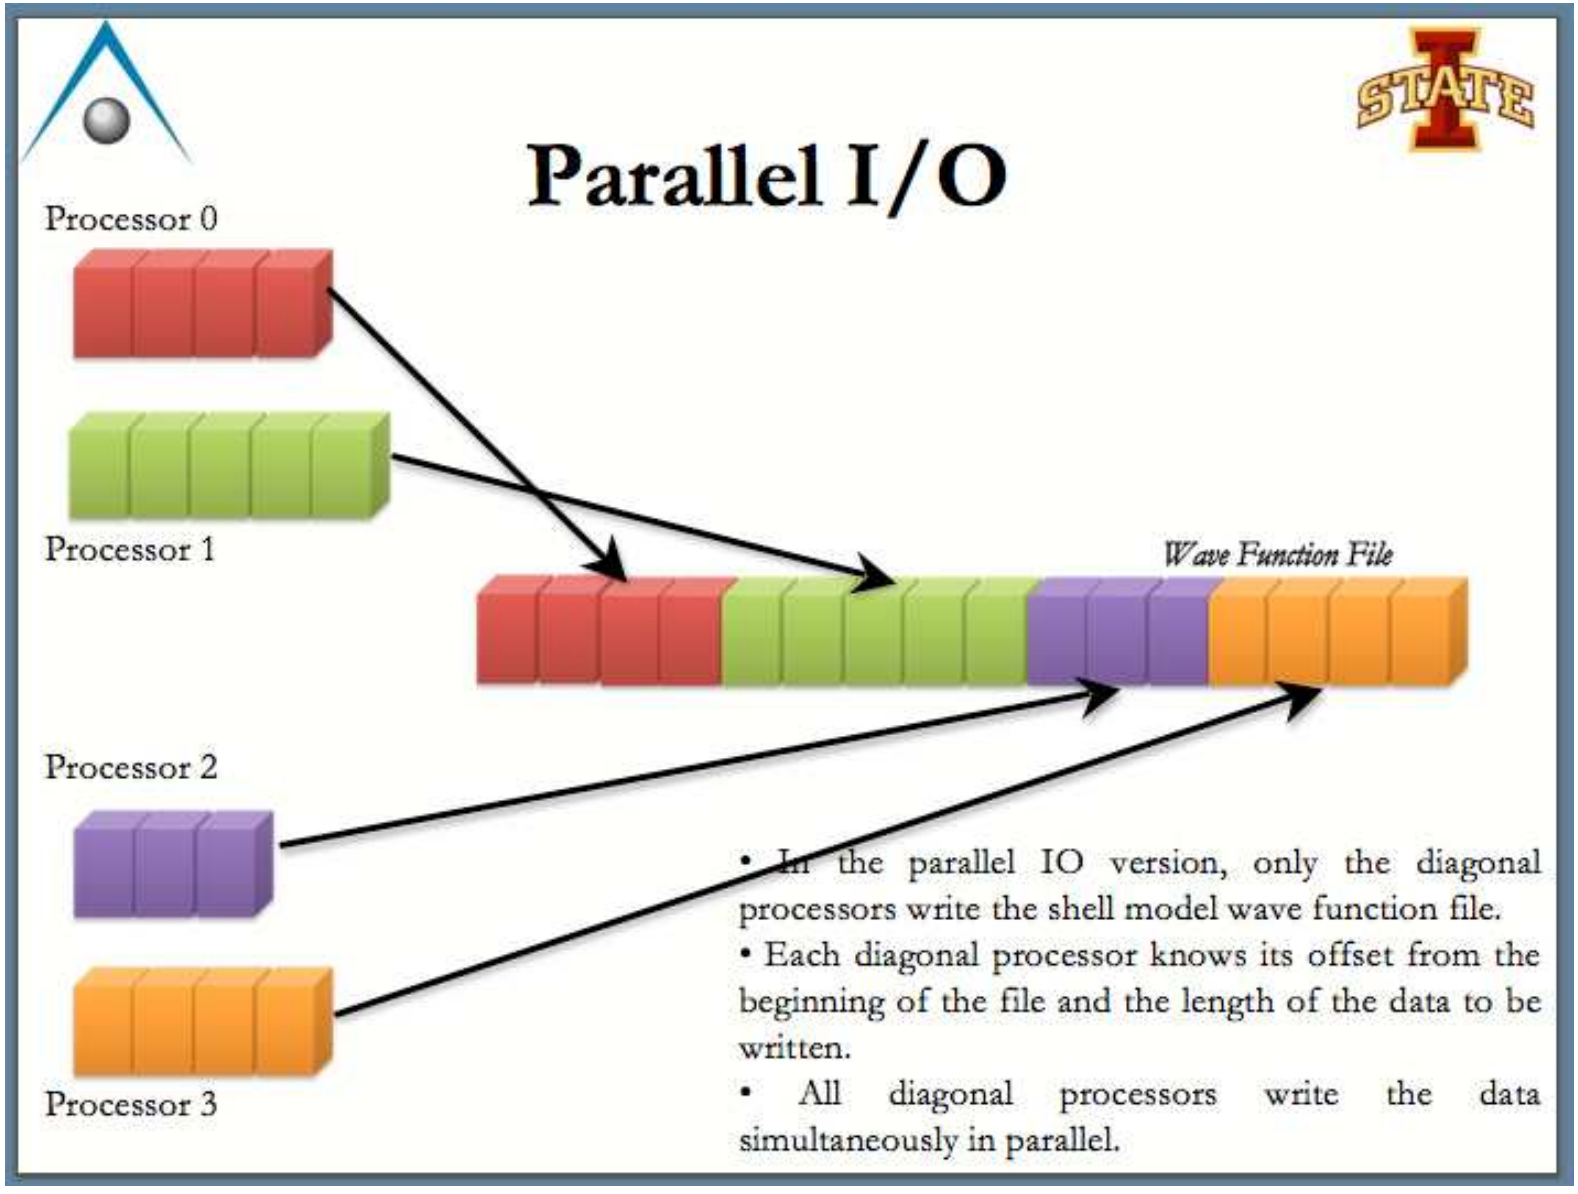

IO of wavefunctions – presented at ICCS 2009

- HDF5: self-describing portable files
- HDF5 more user-friendly than binary
- performance tests on n diagonals on Franklin comparison of sequential binary vs. parallel HDF5

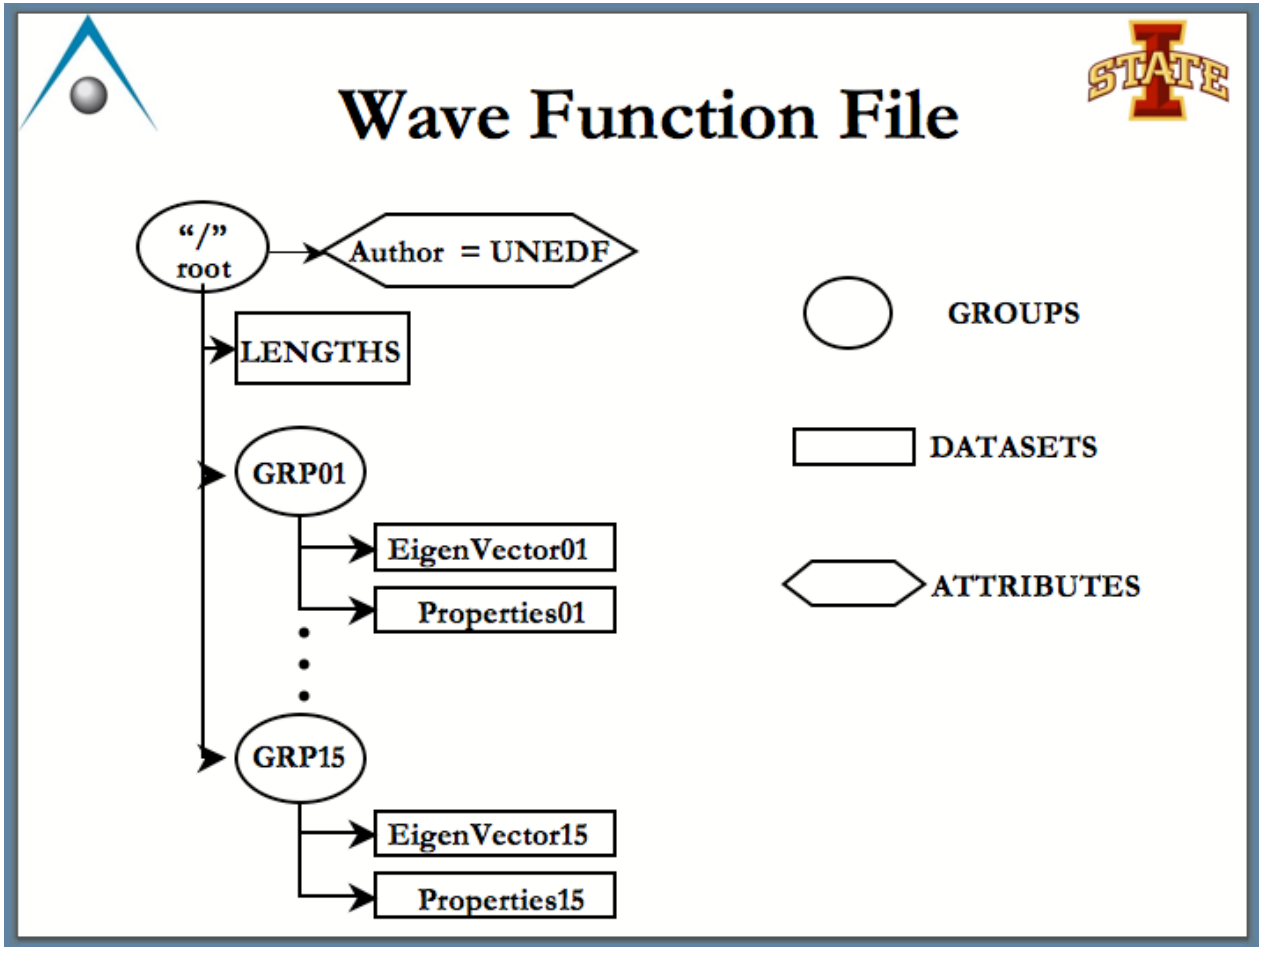

IO of wavefunctions – presented at ICCS 2009

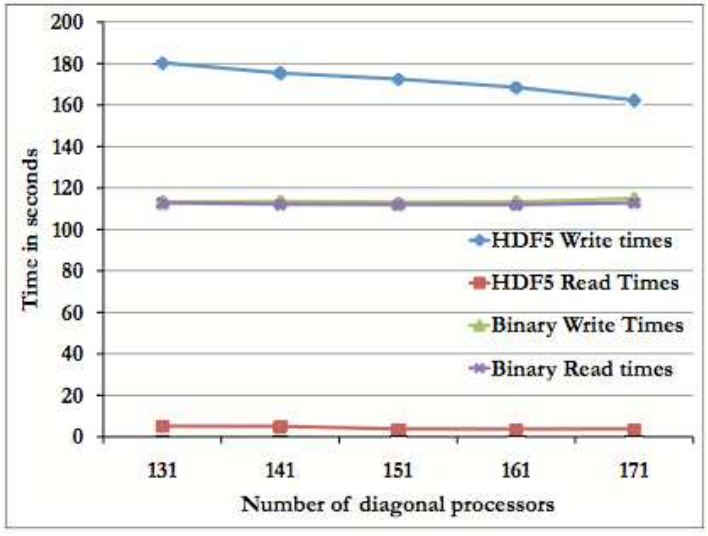

Fig. 4. I/O Performance for 33 GB File

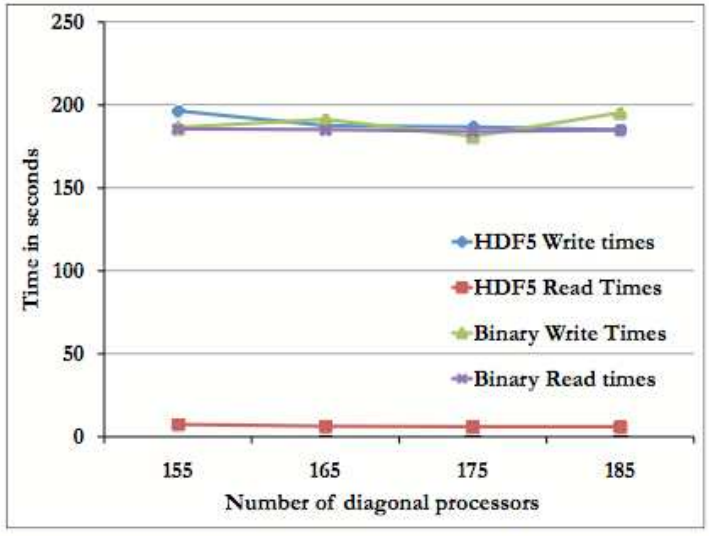

Fig. 5. I/O Performance for 55 GB File

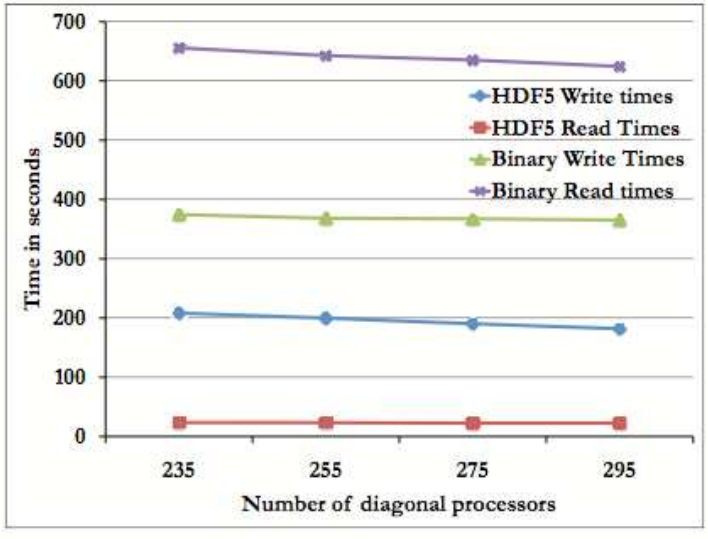

Fig. 6. I/O Performance for 111 GB File

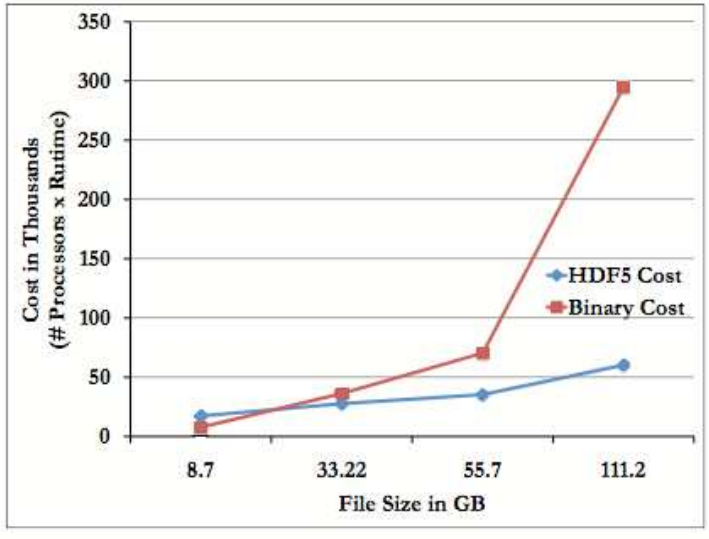

Fig. 7. Cost of using Parallel I/O

IO of wavefunctions – concluding remarks

- Parallel HDF5 more efficient than sequential binary for wavefunction files over 20 GB
- Parallel HDF5 write very slow compared to read
  - probably due to write locks in Lustre file system, no locks for parallel reads
  - PVFS might have better performance, due to absence of write locks?
- Current implementation: MPI_IO
- Parallel HDF5 files more user-friend than (binary) MPI_IO files
- To be done
  - compare efficiency parallel HDF5 vs. MPI_IO

IO Challenges – Checkpointing/restarts

- Many-body matrix: unique for each processor
  - MPI_IO of sparsity index arrays implemented
  - indexing re-usable for subsequent runs of same nucleus in same model space
  - IO of actual matrix elements currently not implemented
  - Challenge:
    - Typical transfer rate JaguarPF (Cray XT5, 150,000 cores): 20 GB/s
    - IO of 220 TB (1.5 GB/core) takes 3 hours . . .

- Lanczos vectors
  - could be written to file after each iteration by “diagonal” processors (currently under consideration)
  - possible use of ADIOS for asynchronous IO? (ORNL)

Work in progress: neutron drops

- Neutrons in H.O. external field
- No extrapolation to infinite model space yet
- No estimates of numerical error bars yet
- Similar runs for odd neutrons n = 3 to 15 in progress

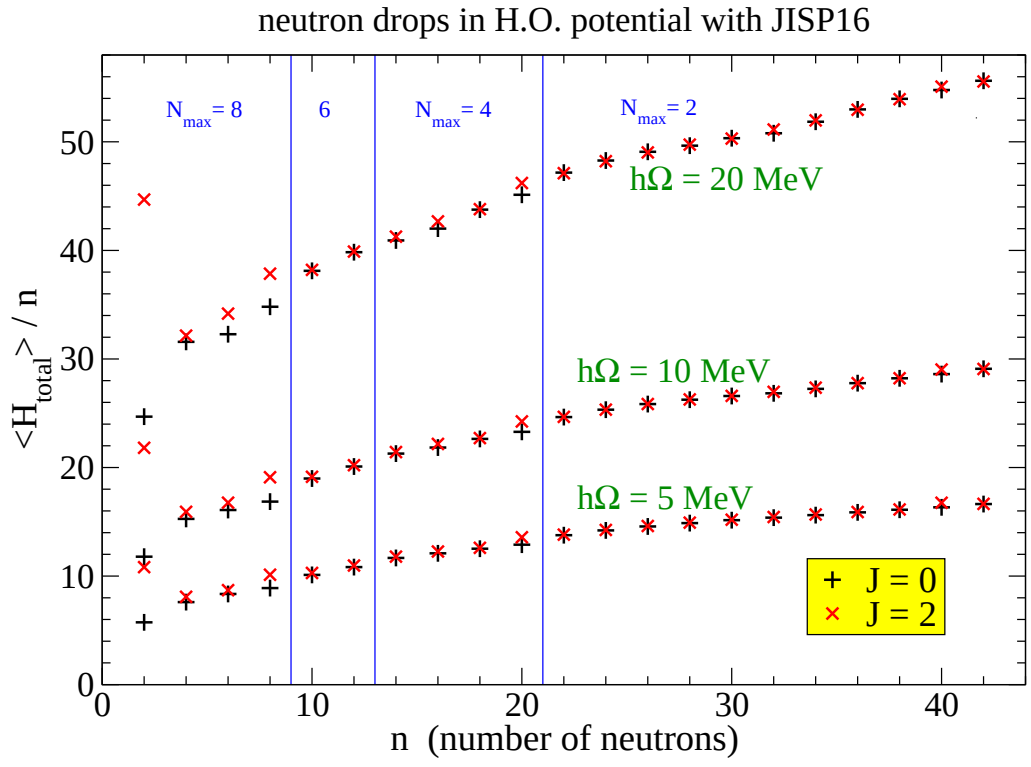

Work in progress: neutron drops

- Neutrons in H.O. external field
- No extrapolation to infinite model space yet
- No estimates of numerical error bars yet
- Similar runs for odd neutrons n = 3 to 15 in progress

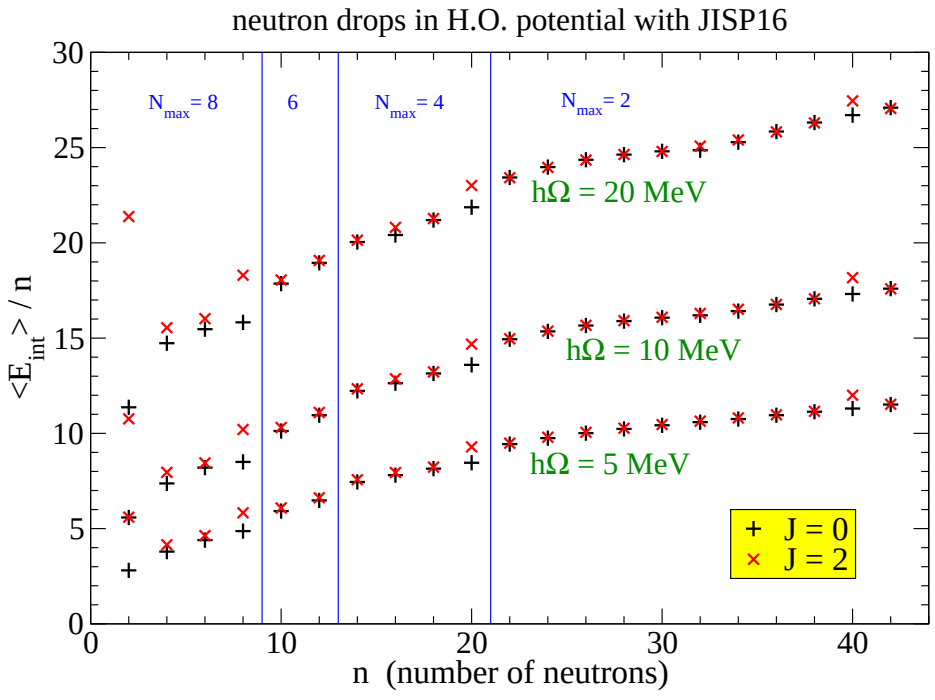

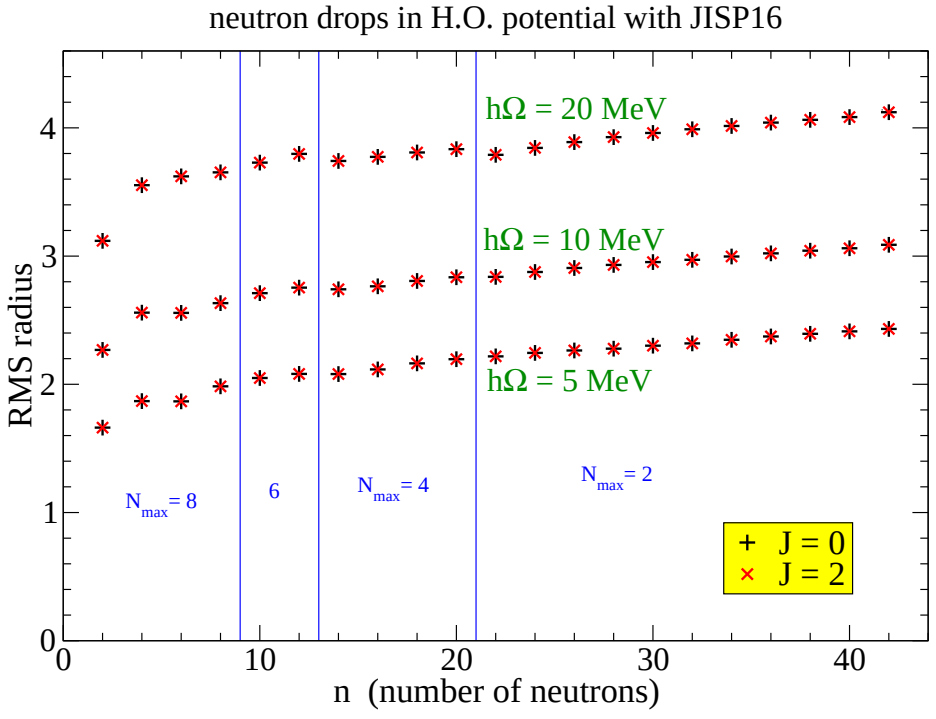

Work in progress: Different truncation schemes

- No-Core Shell Model, No-Core Full Configuration
  - truncation on total number of H.O. quanta, $N_{\max}$, in many-body basis space
- Full Configuration Interaction
  - truncation on single-particle basis space, retaining all many-body states allowed by the symmetries
- Monte-Carlo Shell Model

  Abe, Otsuka, Shimizu, Utsuno
  - sampling of the many-body basis in a FCI truncation
- Importance Sampling

  Navratil, Roth
  - sampling of the many-body basis in a $N_{\max}$ truncation
- Symplectic No-Core Shell Model

Draayer *et.al.*, PetaApps grant

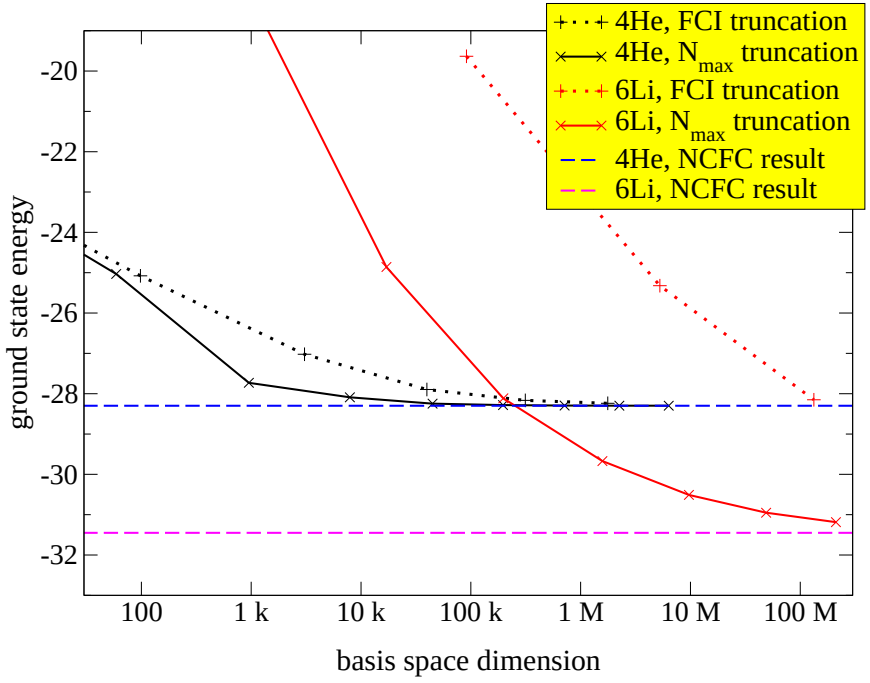

Work in progress: Wood–Saxon basis

- MFDn works with any set of single-particle basis states that can be characterized by quantum numbers $|n, l, j, m_j\rangle$
- MFDn needs N-body interaction in that basis
- Commonly used:
  H.O. basis, which falls off too fast at large distances
- Question:
  Does a more realistic basis exhibit faster convergence, in particular for RMS radii?
  E.g. Wood–Saxon basis

$$U(r) \quad = \quad \frac{U_0}{1 + \exp[(r - R)/a]} + U_{SO}$$

- Need to construct interaction in W.S. basis
- Shortcut used for JISP16:
  Expand W.S. wavefunction in (finite number of) H.O. functions

Initial results with Wood–Saxon basis for 4He

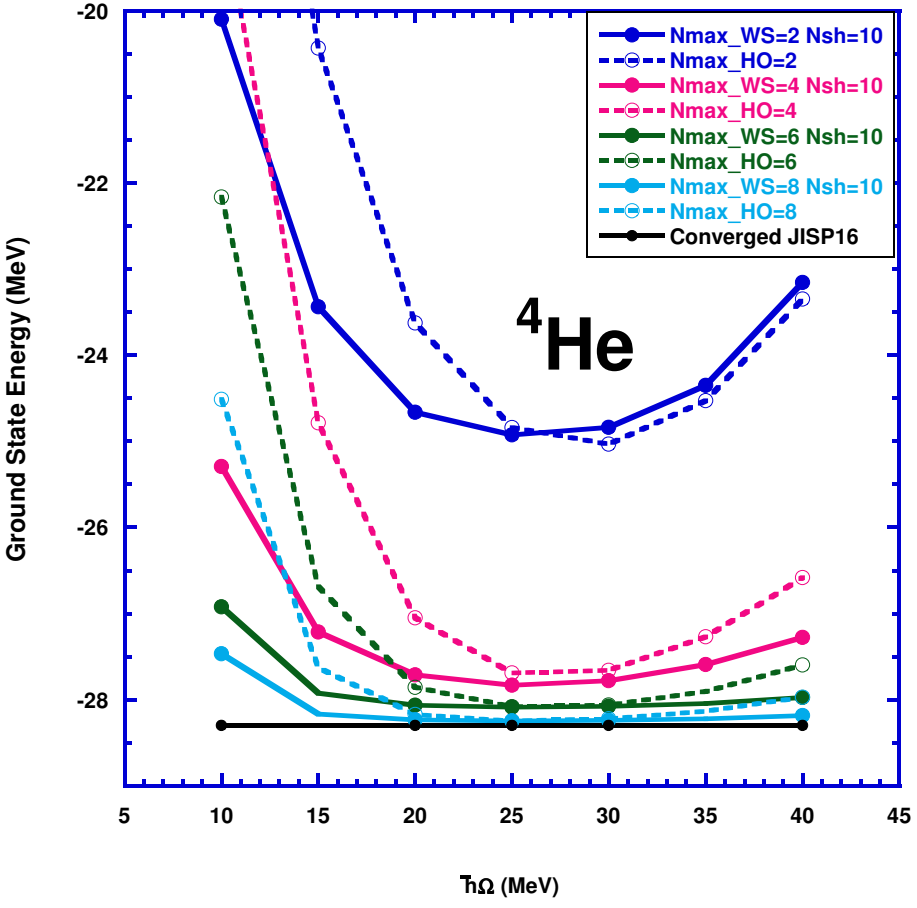

- optimize W.S. parameters by minimizing ground state energy
  - VT-direct (derivative-free): typically 10,000 MFDn runs
  - NEWUOA: typically 100 to 200 MFDn runs

  in order to determine optimal W.S. parameters
- slight improvement of variational bound

Initial results with Wood–Saxon basis for 4He

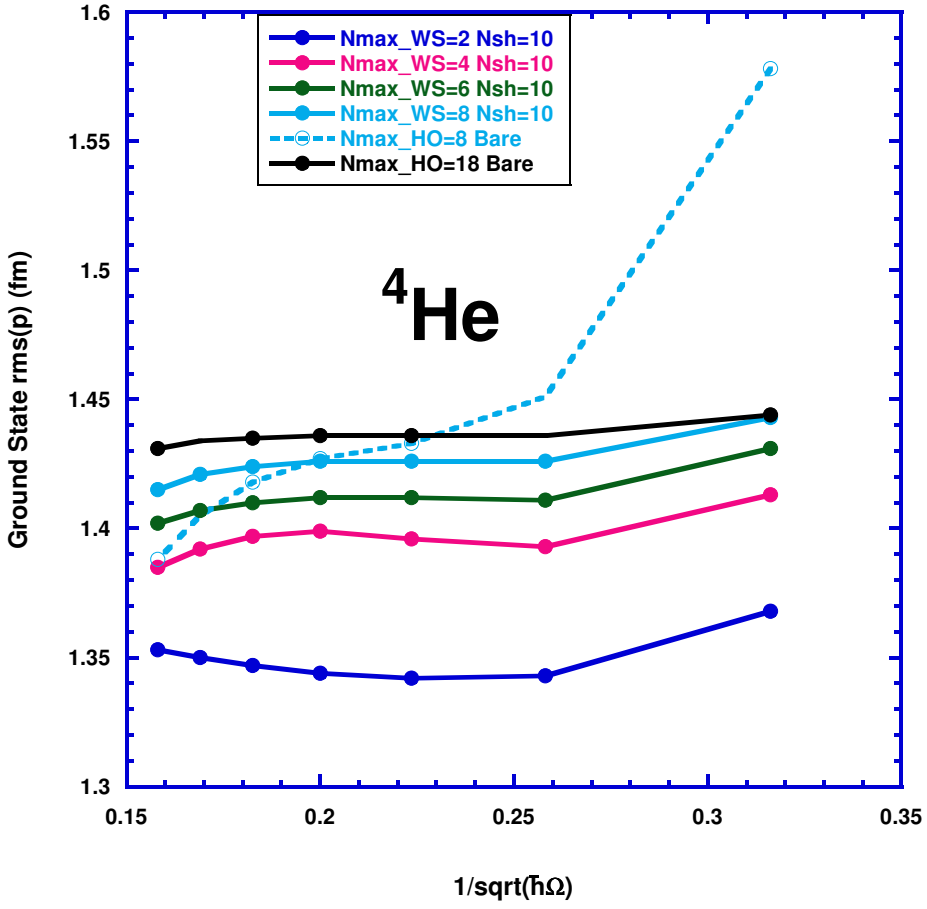

- RMS radius nearly independent of H.O. parameter $\hbar\Omega$ provided that sufficient H.O. levels are used for the expansion of the W.S. basis
- no convergence with $N_{\max}$ yet, need to go to larger model spaces

Initial results with Wood–Saxon basis for 12C

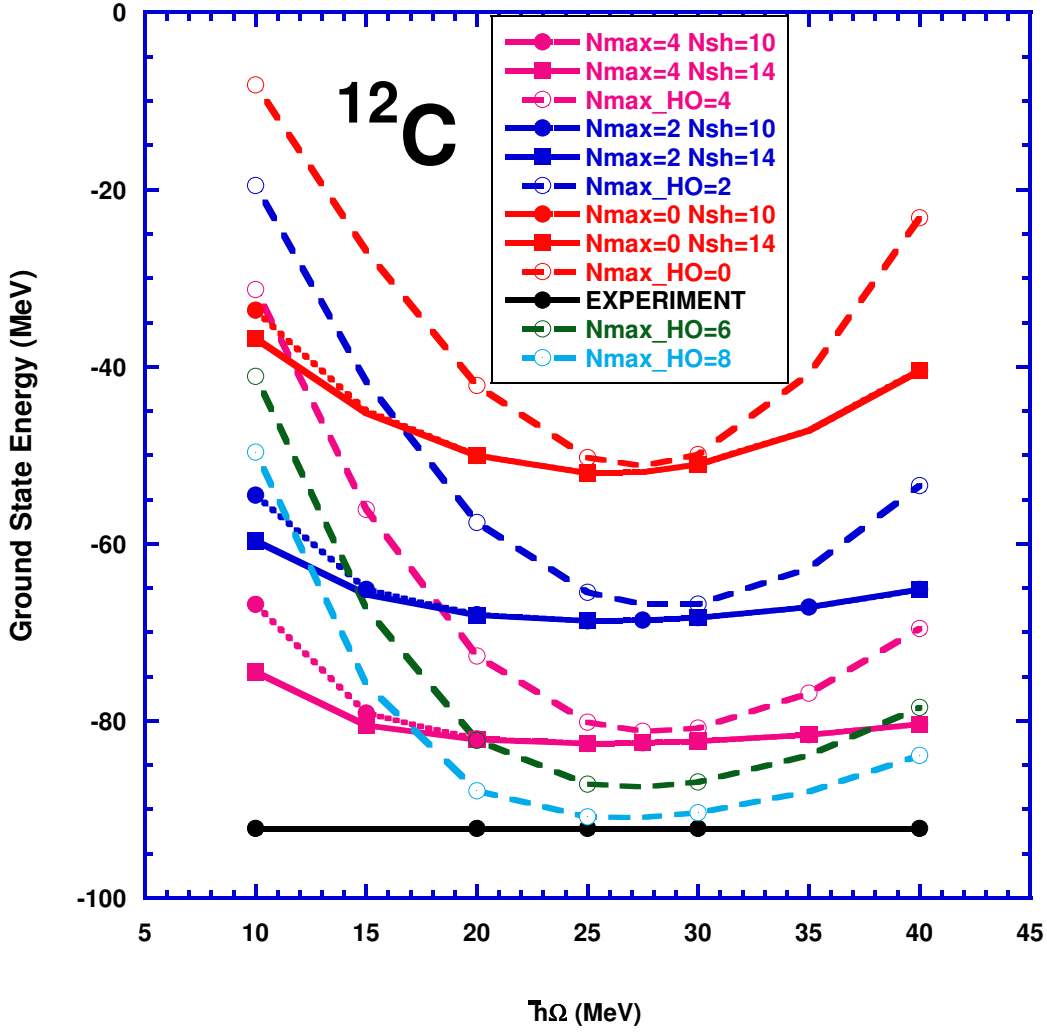

- slight improvement of variational bound with suitable W.S. parameters
- nearly independent of H.O. parameter $\hbar\Omega$ provided that sufficient H.O. levels are used for the expansion of the W.S. basis

Initial results with Wood–Saxon basis for 12C

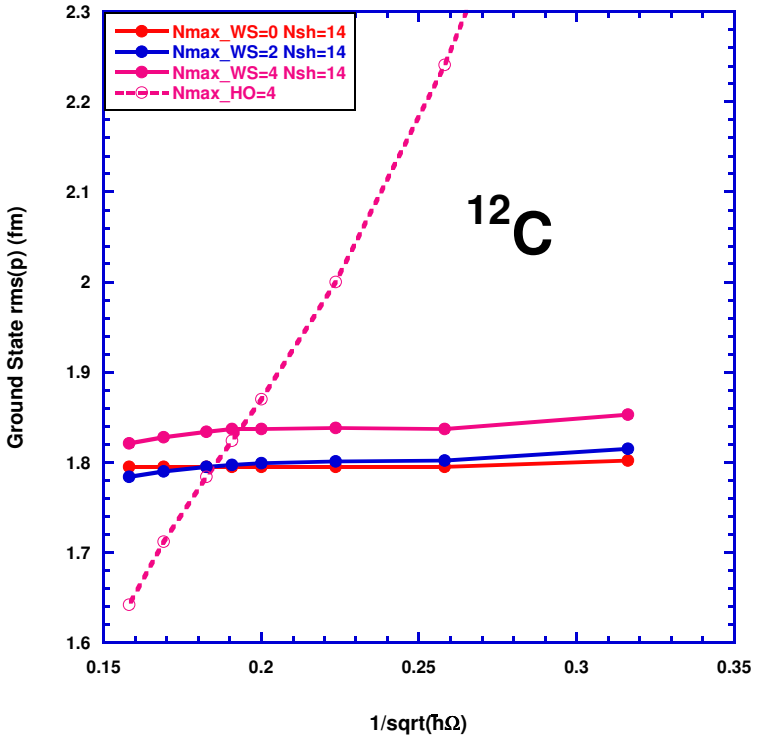

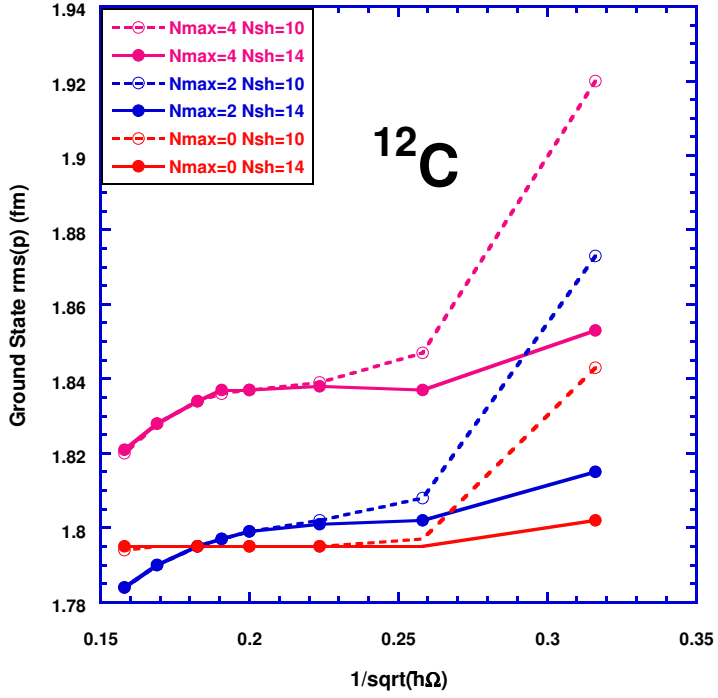

- no convergence with $N_{\max}$ yet, need larger $N_{\max}$ values
- currently studying convergence rate for deuteron, where we can use large model spaces
- future work: He-isotopes

Work in progress: Scattering

J-matrix formalism:

scattering in the oscillator basis

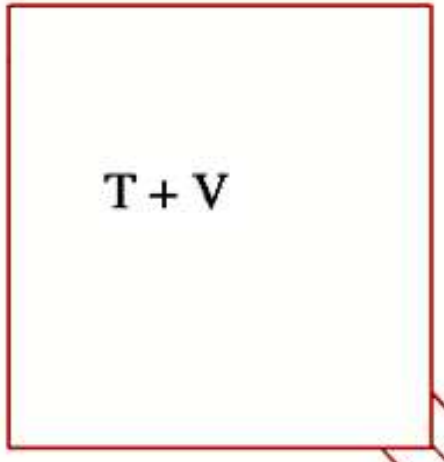

$$\sum_{n'=0}^{N} H^{l}_{nn'} \langle n' | \lambda \rangle = E_{\lambda} \langle n | \lambda \rangle, \qquad n \le N$$

$$G_{NN}(E) = -\sum_{\lambda=0}^{N} \frac{\langle N | \lambda \rangle^2}{E_{\lambda} - E}$$

$$S = \frac{C^{(-)}_{Nl}(q) - G_{NN}(E) T^{l}_{N,N+1} C^{(-)}_{N+1,l}(q)}{C^{(+)}_{Nl}(q) - G_{NN}(E) T^{l}_{N,N+1} C^{(+)}_{N+1,l}(q)}$$

n(p)+nucleus applications

Forward scattering J-matrix
1. Calculate E_{λ} and $\langle N | \lambda \rangle$ with NCSM(JISP16)
2. Solve for S-matrix and obtain phase shifts

Inverse scattering J-matrix
1. Obtain phase shifts from scattering data
2. Solve for n(p)+nucleus potential, resonance params

T

A.M. Shirokov, A.I. Mazur, J.P. Vary, and E.A. Mazur, arXiv:0806.4018; and references therein

Work in progress: Scattering

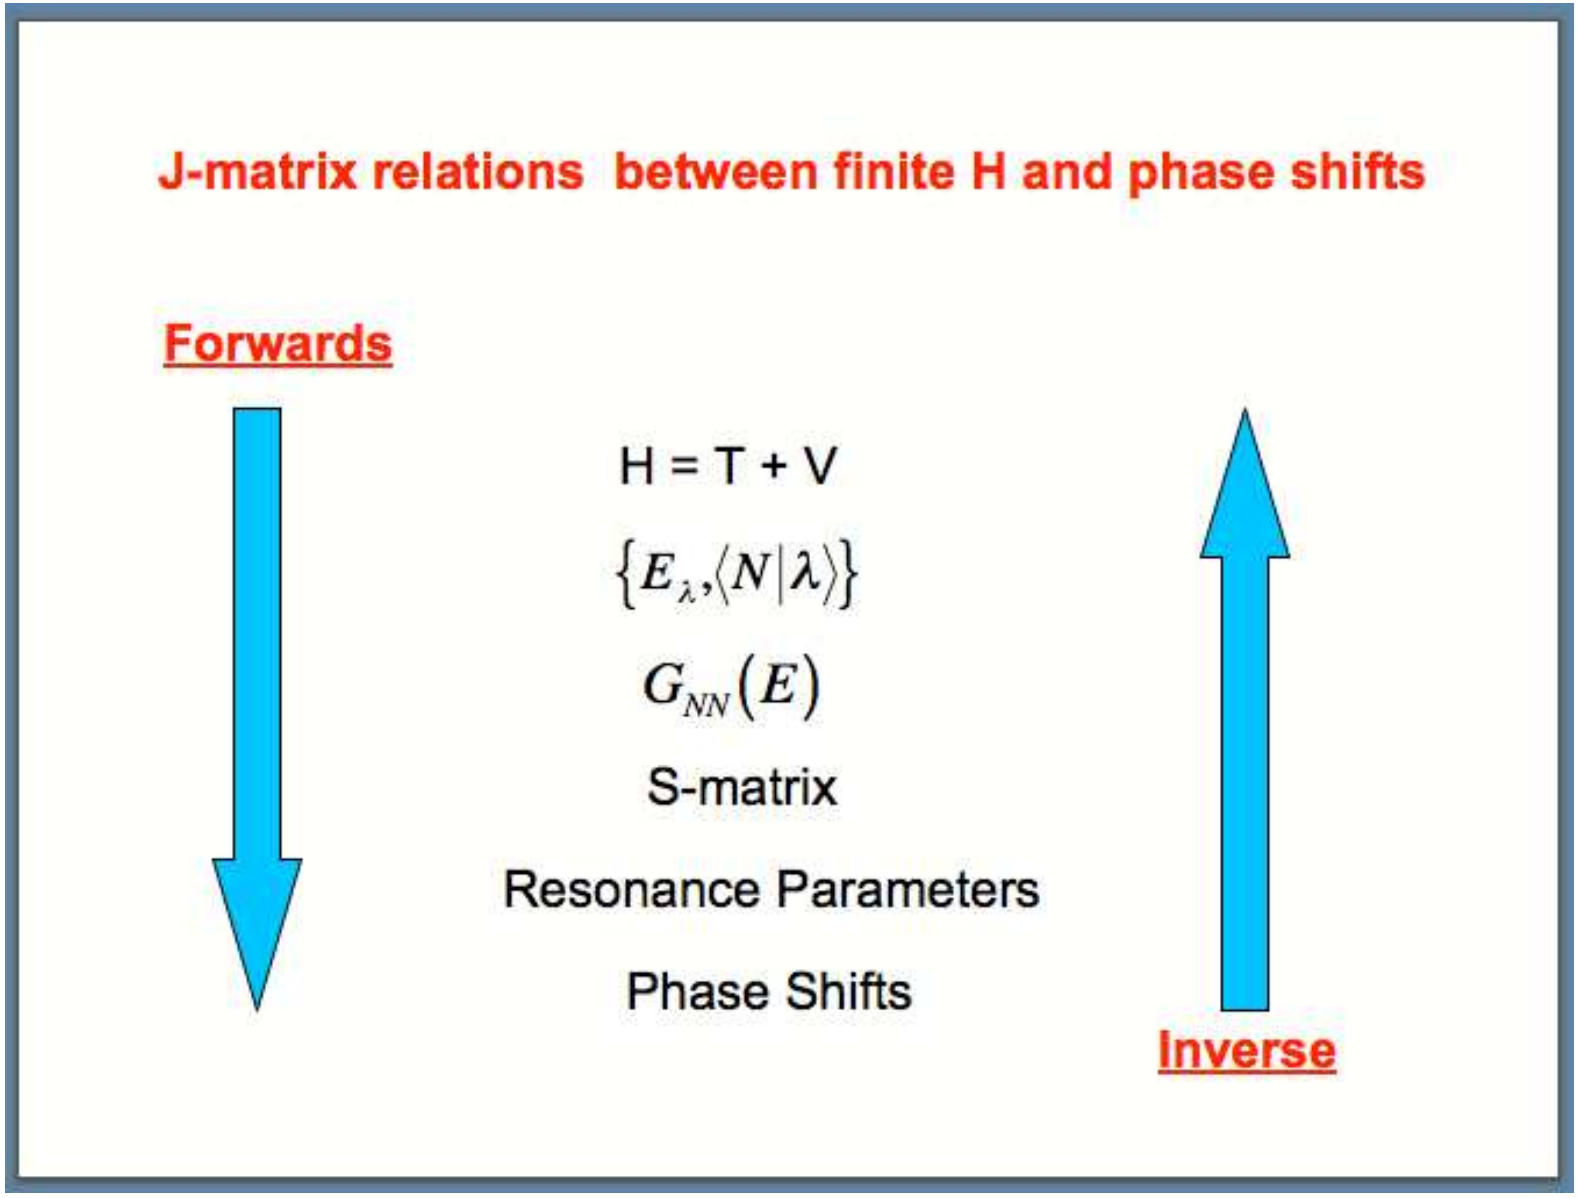

Work in progress: $n\alpha$ Scattering with JISP16

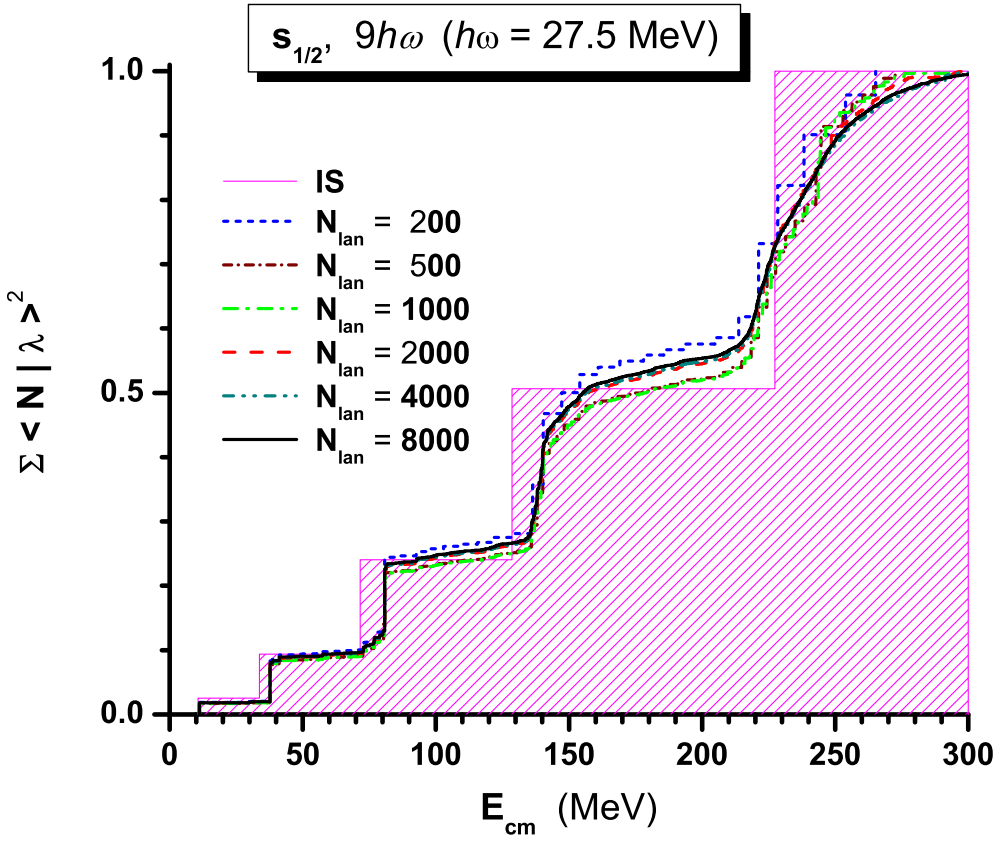

$\langle \mathcal{N} | \lambda \rangle$ is component of ^{5}He wavefunction with α core and one neutron in highest orbital

- Inverse Scattering: can fit phase shifts with 5 scattering states
- NCSM: in principle all excited states that are free of Center-of-Mass motion

NCSM and IS give qualitatively similar results for $\sum \langle \mathcal{N} | \lambda \rangle^2$ as function of E, provided one incorporates “all” excited states

Work in progress: $n\alpha$ Scattering with JISP16

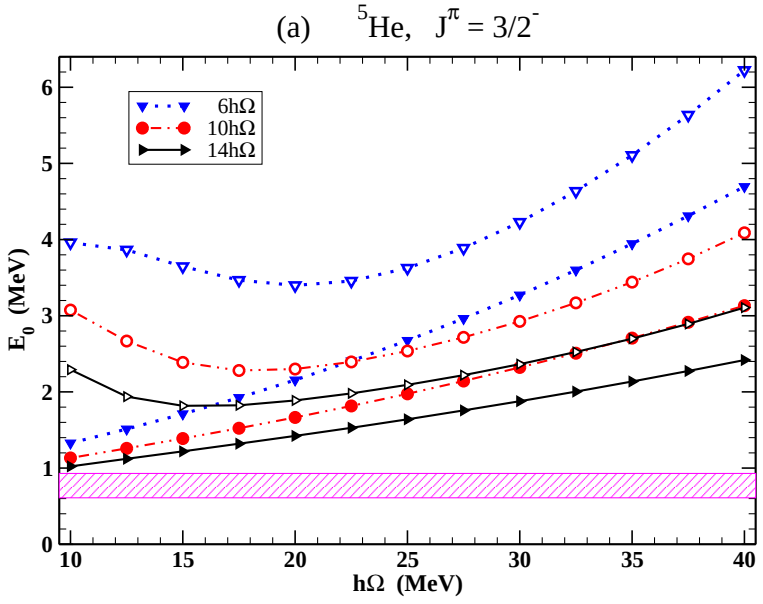

- Inverse Scattering: E_λ depends almost linearly on H.O. basis parameter $\hbar\Omega$
- NCSM calculations for resonance states exhibit similar $\hbar\Omega$ dependence

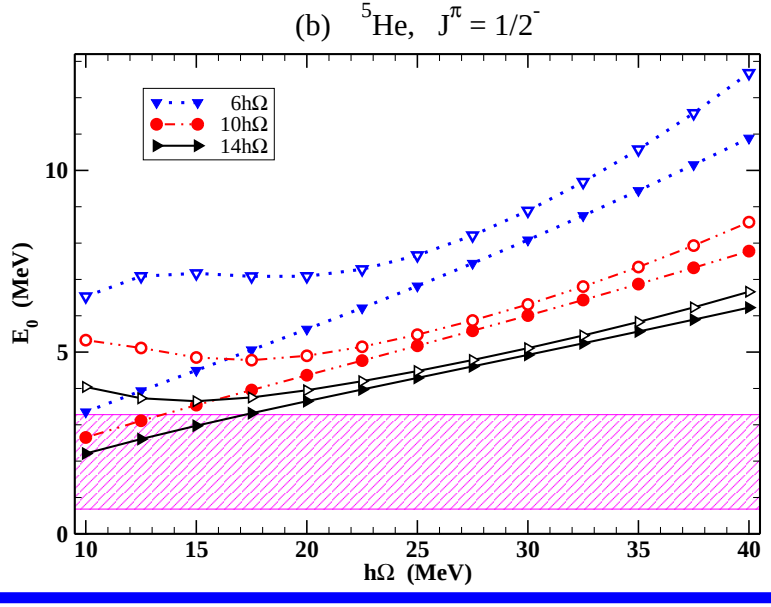

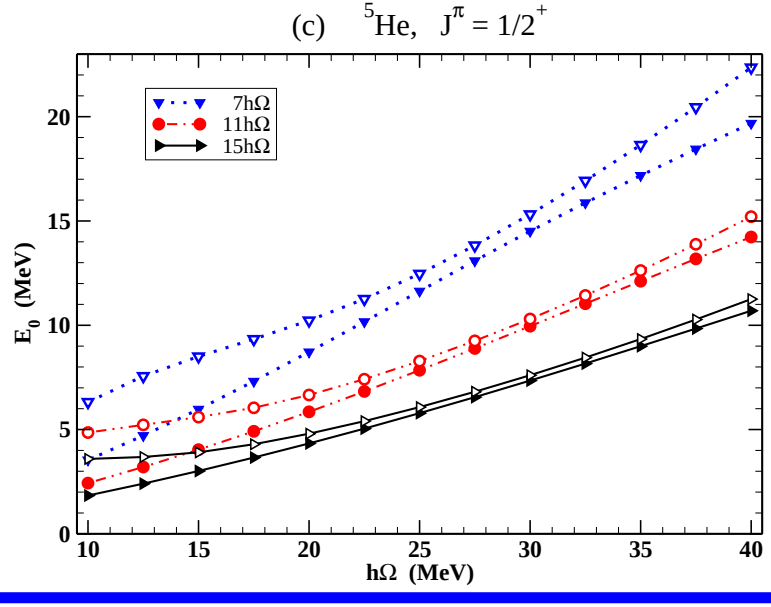

Work in progress: $n\alpha$ Scattering with JISP16

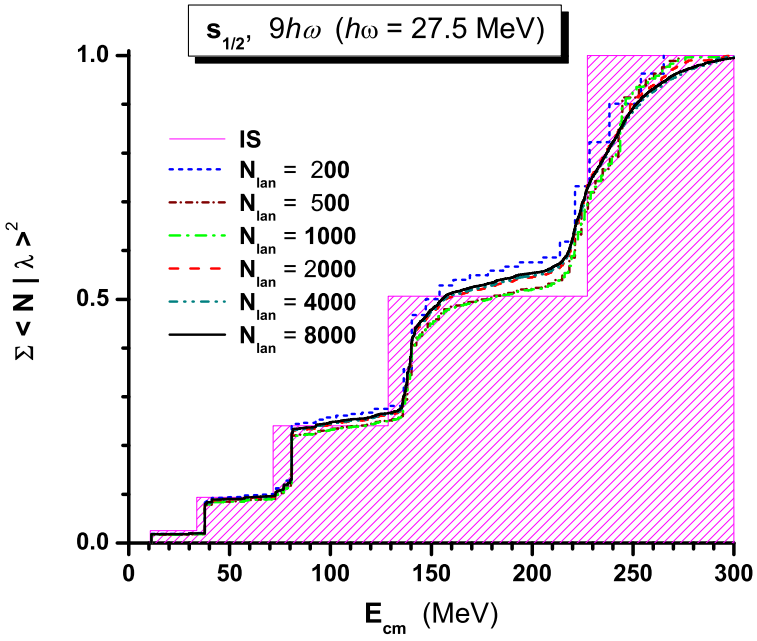

Challenges

- separation of CoM excitations for highly excited states
- construction of phase shifts from $\mathcal{G}_{\mathcal{N}\mathcal{N}}(E)$ very sensitive to E_λ and $\langle \mathcal{N}|\lambda\rangle$
- slow convergence with $N_{\max}$, especially at larger energies

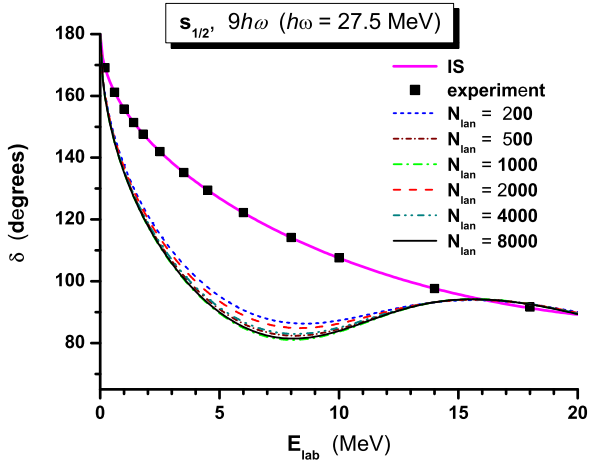

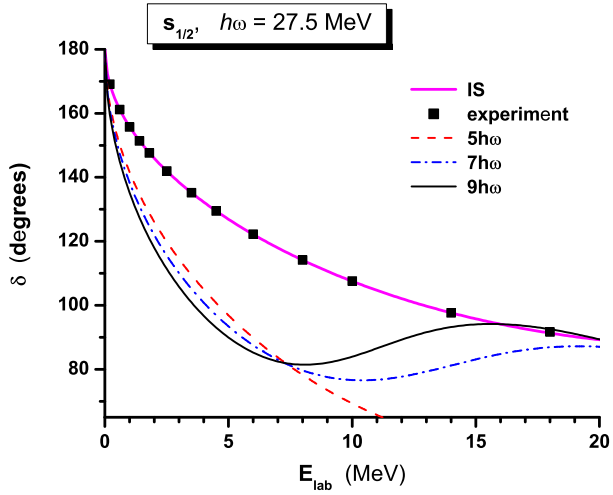

Work in progress: Total-J code

Under development

- Total-J basis:
  eigenstates of both $\hat{M}$ and $\hat{J}^2$
  - significant reduction of
    basis space dimension D

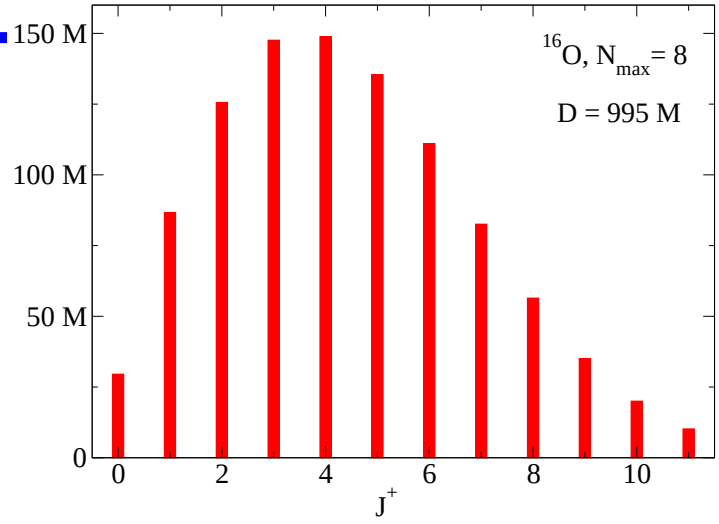

- Key observation: $\hat{J}^2$ changes quantum number m_i only, but not other quantum numbers (n_i, l_i , j_i) of a single particle state
  - arrange many-body states in groups
    with each nucleon in a specific (n, l , j) orbital
  - diagonalize $\hat{J}^2$ within each group and use eigenvectors with specific eigenvalues $J(J+1)$ to construct transformation matrices Q from m-scheme basis to Total-J basis
  - construct many-body matrix in Total-J basis $H^J = QH^mQ^T$
  - better yet, use QR factorization ? (*to be implemented*)

Work in progress: Total-J code

Preliminary results (single processor program)

- Matrix H^J has dense-block structure
  - more efficient MATVEC's using standard libraries
- Can obtain large number of states with specific spin only
  - for scattering of neutron on ^{4}He,
    we need "all" $J = \frac{1}{2}$ and $J = \frac{3}{2}$ states of ^{5}He

$N_{\max} = 8$	dimension	# nonzero m.e.	CPU time	
m-scheme	271,023	114,772,843	2 hrs (15 pe's)	100 states 22 $J = \frac{1}{2}$ 31 $J = \frac{3}{2}$
$J = \frac{1}{2}$	24,804	21,775,381	18 min	100 states
$J = \frac{3}{2}$	43,211	66,152,353	31 min	100 states

  - for E(2) strength function of spin-0 nuclei,
    need "all" $J = 2$ excited states (quadrupole excitations)

Work in progress: Total-J code – Challenges

- Groups of many-body states w. nucleons in specific (n, l , j) orbitals can contain 10^5 (or more?) of many-body states
  - need to diagonalize thousands (millions?) of matrices with dimensions of up to 10^5 (or even bigger?) in order to set up our Total-J basis space. . .
  - load balancing and scalability ?
- Total-J basis is smaller than m-scheme dimension, but matrix H^J can have (significantly) more nonzero elements then H^m

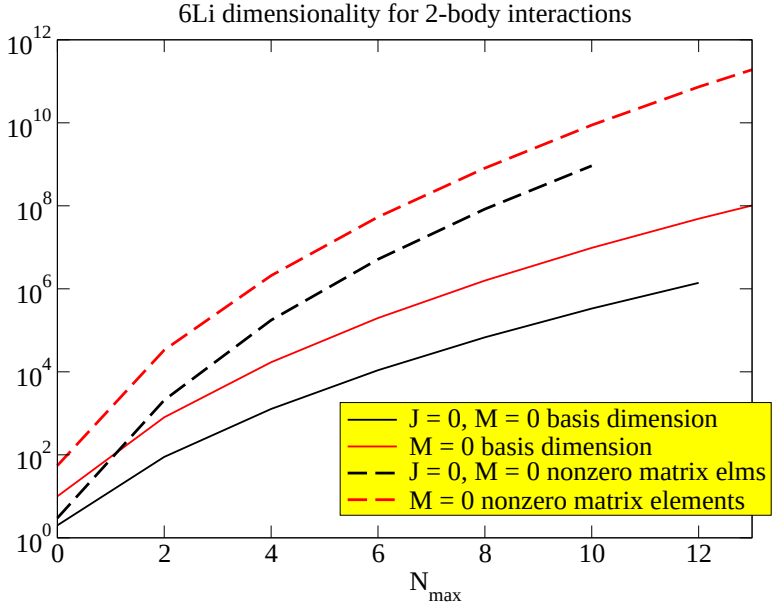

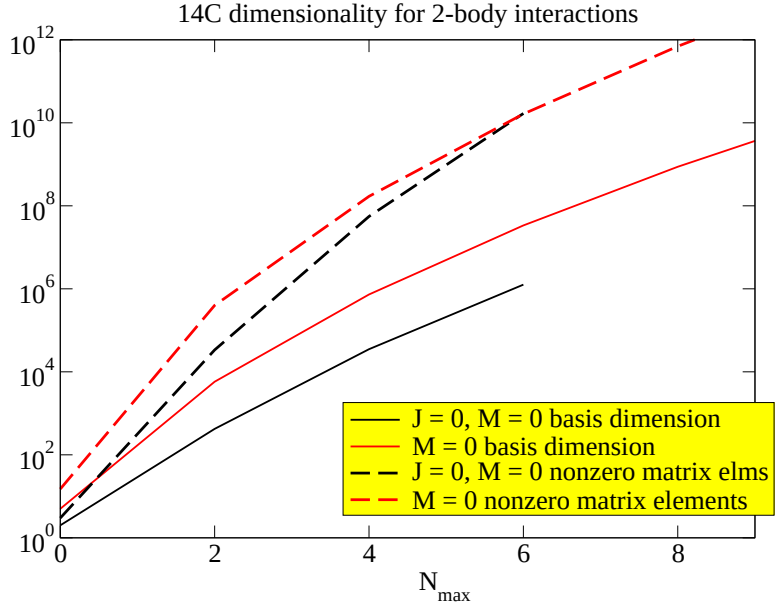

Work in progress: Be-isotopes

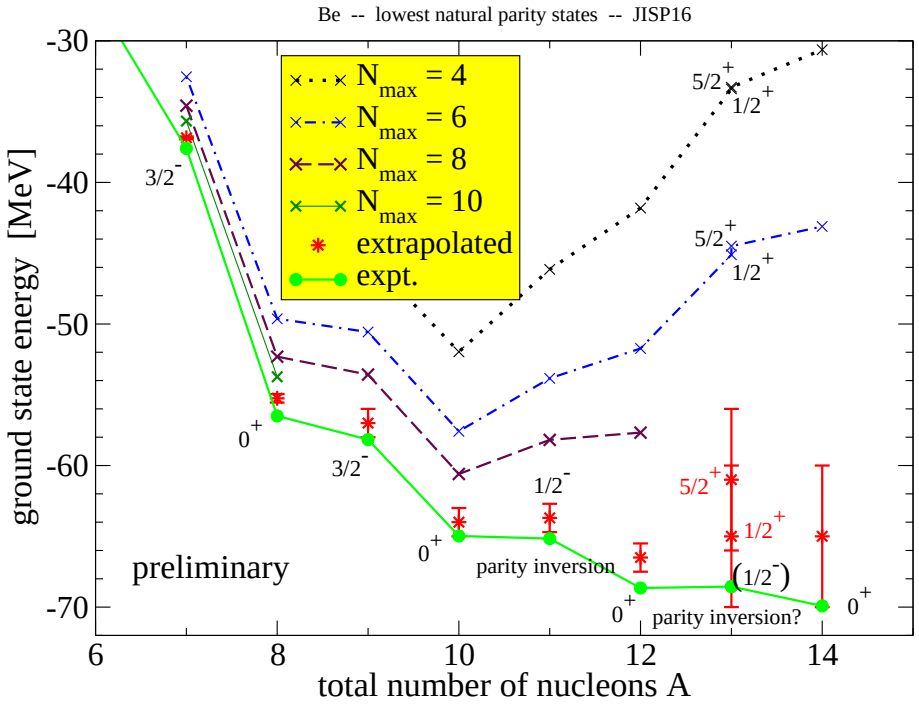

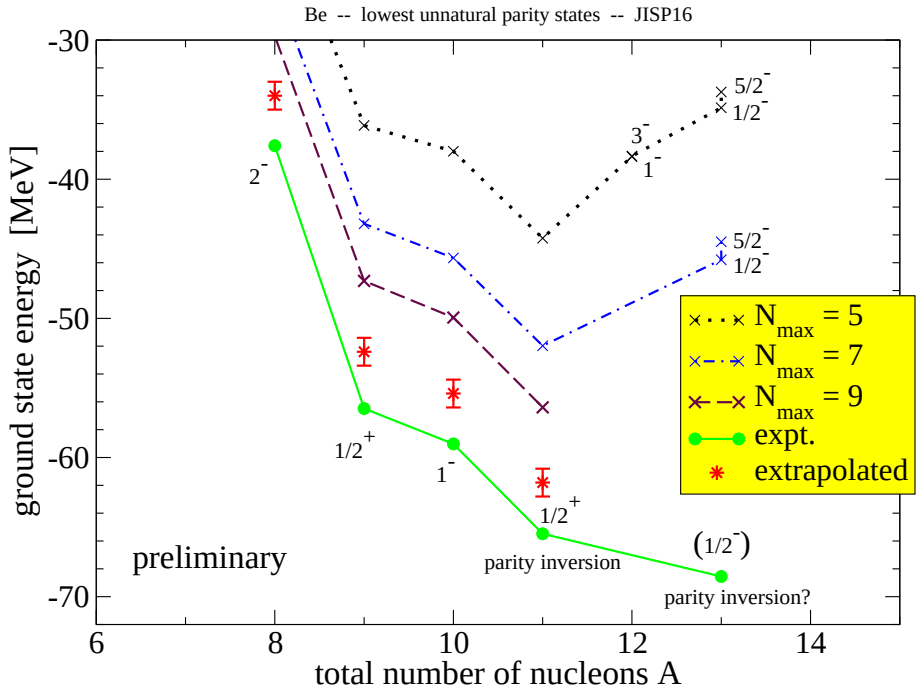

- Exploring physics near the neutron drip line in progress
- Similar results for He- and Li-isotopes

Petascale Early Science – Ab-initio structure of Carbon-14

- 30,000,000 CPU hours, one of 12 awards on JaguarPF
- Collaborators: David Dean, Hai Ah Nam (ORNL), Petr Navratil, Erich Ormand (LLNL), James Vary, PM (ISU)

Puzzling to scientists …

What is the nuclear structure of ^{14}C that leads to its anomalously long half-life?

$\tau_{1/2}$ = 5730 years

^{10}Be and ^{14}C have extremely long half-lives compared to other light nuclei (1.6 x 10^6 years / 5,730 years) . Their long half-lives make both isotopes useful for radioactive dating.

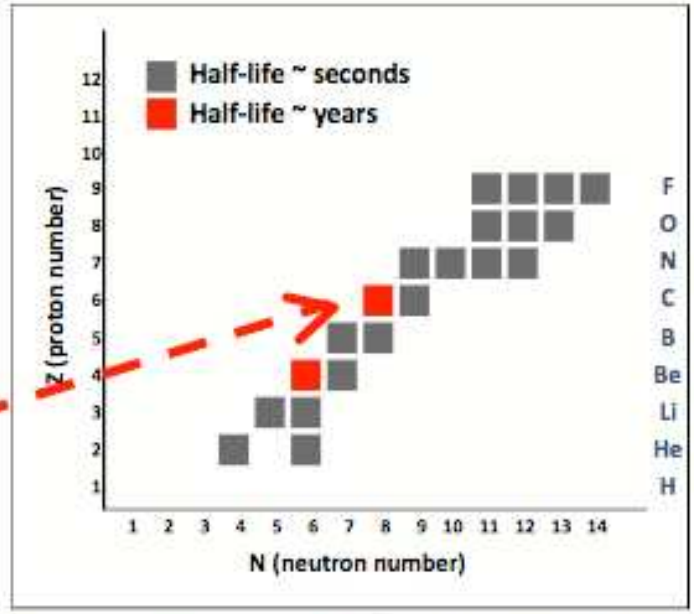

Chart of light nuclei that decay via beta emissions

$n \rightarrow p + e^- + \nu_e$

Petascale Early Science – Ab-initio structure of Carbon-14

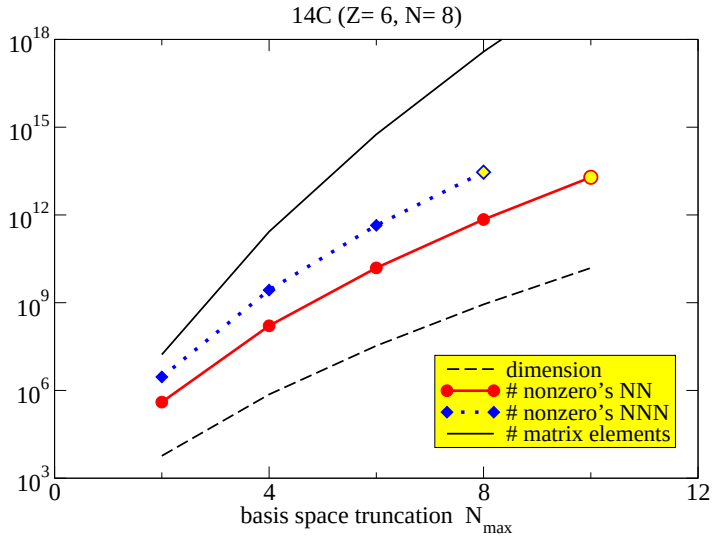

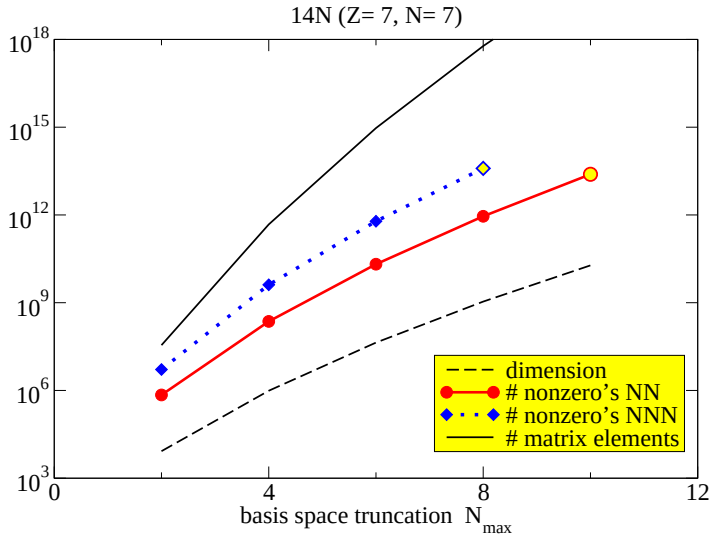

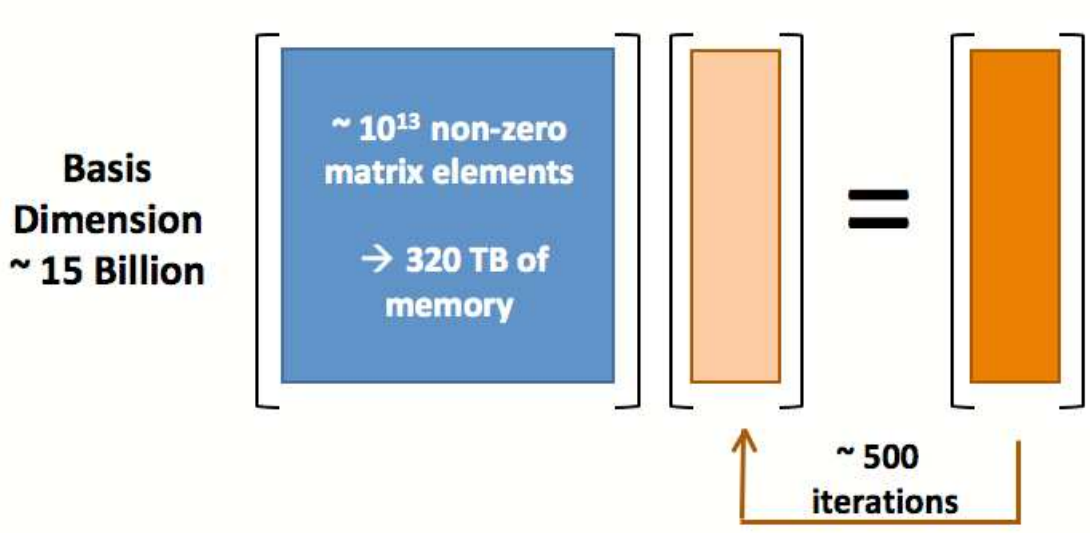

might fit on 150k cores with some additional code-development . . .

Petascale Early Science – current status

- code is memory bound
  - partial “out-of-core” calculations ? too slow
  - compression of indexing arrays ? about 10% savings
  - using Co-Array Fortran for 3-body interaction ? too slow
  - partial “regenerate-on-the-fly” ? promising !
    14C, chiral 3-body forces, $N_{\max} = 8$, on 110,920 cores
    - about 1.5 hr for setting up many-body matrix arrays
    - about 3 minutes per Lanczos iteration
- I/O of Lanczos vectors for checkpointing/restarts implemented
- further improvements needed
  - single-processor performance of regenerate-on-the-fly code
  - load balancing of regenerate-on-the-fly code
  - hybrid MPI/OpenMP performance
  - re-asses Co-Array Fortran after interconnect upgrade
    (at the end of this calendar year)

Progress report – Year 3

- Multilevel blocking – automated blocking installed
- Multi-threading – done, performance depends on application
- Total-J – serial code works, parallel code to be developed
- Improved IO – on track
- Improved chiral interactions
  - superseeded by LLNL investigations of 3-body β decay
  - runs for A=7, 8, 12, and 14 in progress
- Expanded use of realistic basis functions
  - MFDn interfaced with VTdirect and NEWUOA to facilitate automatic searches for optimal basis space parameters
- Extension of ab-initio approach to scattering – on track
- Expanded use of external fields – on track
  runs for $n = 6$ to 20 neutrons with JISP16 in progress
- One-body density matrix – to be done

- Evaluate χNNN matrix on-the-fly – waiting on collaborators

Roadmap – Year 4

- Develop efficient algorithms to compute null-space of $\hat{J}^2$ sub-matrices (LBNL/ISU)
- Improve IO for checkpoints/restarts (AL/ISU/ORNL)
- Investigate alternative eigensolver, especially the advantages of block algorithms (LBNL/ISU)
- Improve single-processor performance and load balancing of partial “regenerate-on-the-fly” (LBNL/ISU)
- Set of neutron properties for DFT/DME communities
  - JISP16 for $n = 6$ to 20 aimed at 1% precision (ISU)
  - 2- and 3-body N3LO forces ? (ISU/LLNL?)
- Chiral 2- and 3-body runs for A = 9 through 16 (ISU/LLNL)
- Continue investigation of realistic basis functions (ISU)
- Continue scattering inversigations (ISU/OU)
- Investigations of nuclear matter neutron matter (ISU/MSU)