Large–scale No-Core Shell Model / No-Core Full Configuration calculations

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
Markus Kortelainen, Mario Stoistov, Witek Nazarewicz
OSU: Dick Furnstahl, students
MSU: Scott Bogner
WMU: Mihai Horoi
ANL: Harry Lee, Steve Pieper, Bob Wiringa
LANL: Joe Carlson, Stefano Gandolfi
UA: Sid Coon, Bira van Kolck, Bruce Barrett
LSU: Jerry Draayer, students/postdocs

International Collaborators:
Shirokov, Mazur, Forssen,
Abe, Otsuka, Shimizu, Utsuno

Computer Science/Applied Math

Ames Lab: Masha Sosonkina,
Fang (Cherry) Liu, students
LBNL: Esmond Ng, Chao Yang,
Chris Calderon
ANL: Stefan Wild
OSU: Umit Catalyurek

No-Core Shell Model / No-Core Full Configuration calculations

Numerical approach for bound states of nucleons
using basis-space expansion methods (CI 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
  - selected observables:
    rms radii, magnetic moments, quadrupole moments,
    transition rates between states within the same nucleus
- Ab initio calculations for nuclei throughout the p-shell
  and into the sd-shell with realistic NN and NNN potentials
- Ab initio calculations for nuclei and neutron droplets
  in external fields for comparisons with DME/DFT

Many-Fermion-Dynamics for nuclear physics

- Platform independent hybrid OpenMP/MPI Fortran code
- $N_{\max}$ truncation and HO basis:
  exact factorization CoM motion and intrinsic motion
- No-Core Shell Model:
  improved convergence in relatively small model spaces w. Lee–Suzuki–Okamoto renorm. truncated at 2- or 3-body level
- No-Core Full Configuration:
  monotonic approach to asymptotic values with increasing basis
  - Variational: upper bound for the ground state energy for any finite truncation of the basis space
- Convergence: observables independent of $N_{\max}$ and $\hbar\omega$
- Same interaction, different methods (CC, GFMC, NCFC, NCSM, . . .) give same results within numerical errors

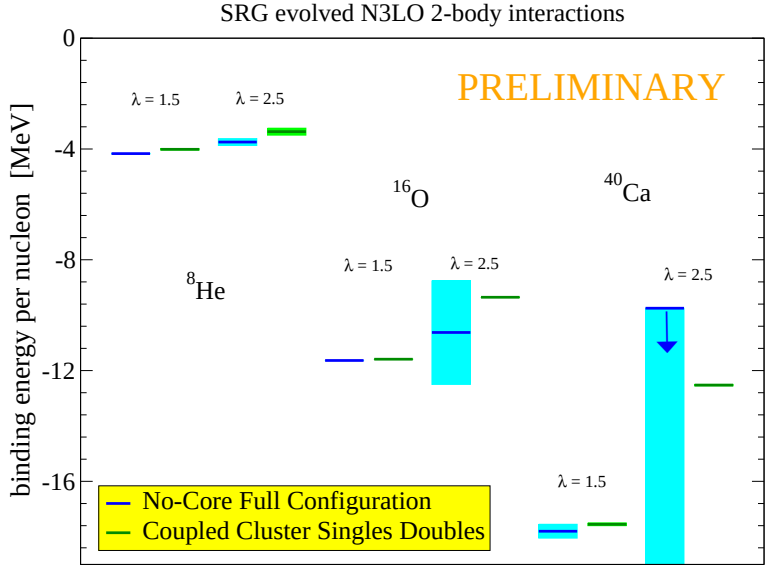

CI calculations – main challenges

- Single most important computational issue:
  exponential increase of dimensionality with increasing HO levels
- Additional computational issue:
  sparseness of matrix / number of nonzero matrix elements

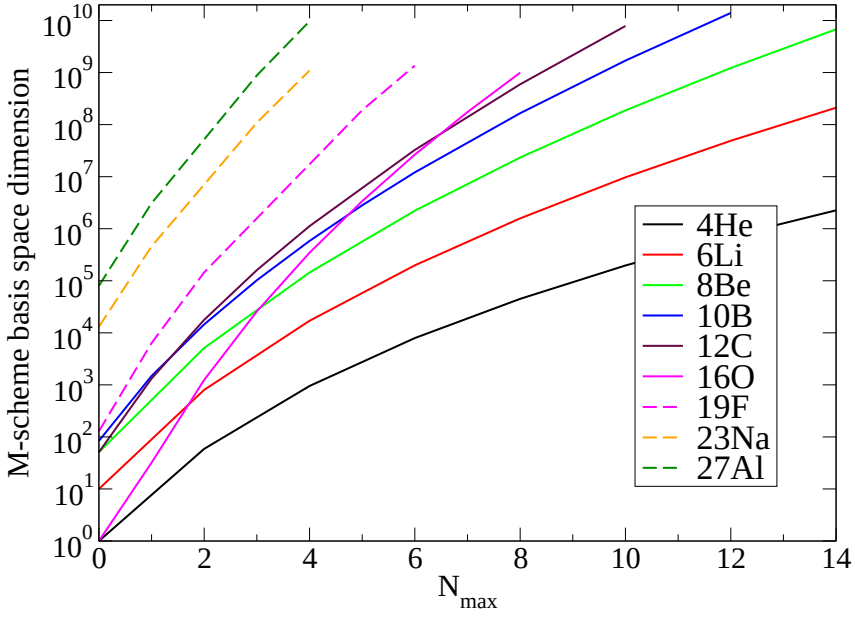

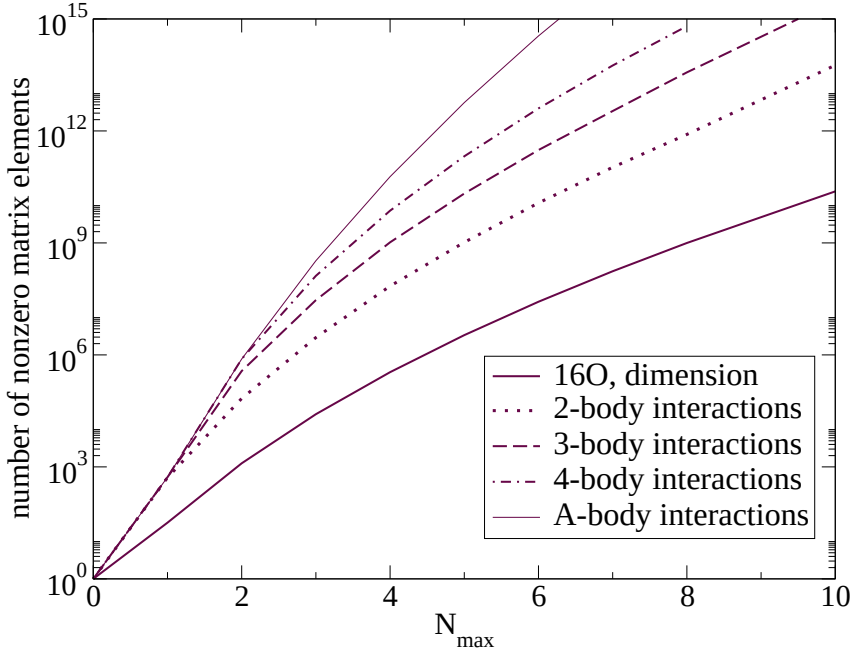

Code development progress report

- MFDn Version 13

  `NERSC:/project/projectdirs/unedf/lcci/MFDn/`

  - improved load-balancing of hybrid OpenMP/MPI code
  - evaluation of One-Body Density Matrix Elements in MFDn
  - developed interface between MFDn and NEWUOA
  - restored interface with TRDENS
  - developing prototype database for archiving ab initio Shell Model / Configuration Interaction results see LCCI session
  - developing interactive python script for running CI codes on Leadership Class facilities see LCCI session

- New: MFDn Version 14 (under development)
  - different distribution of basis states
  - enabling larger model-space calc. using partial “on-the-fly”
  - checkpoint / restart capability

- New: Parallel Total-J code (under development) see next talk

MFDn – Load-balancing

- Each processor has almost the same number of basis states
- Setup time of matrix proportional to number of matrix elements
- Matvec workload proportional to number of matrix elements
- Round-robin distribution of basis states over d procs

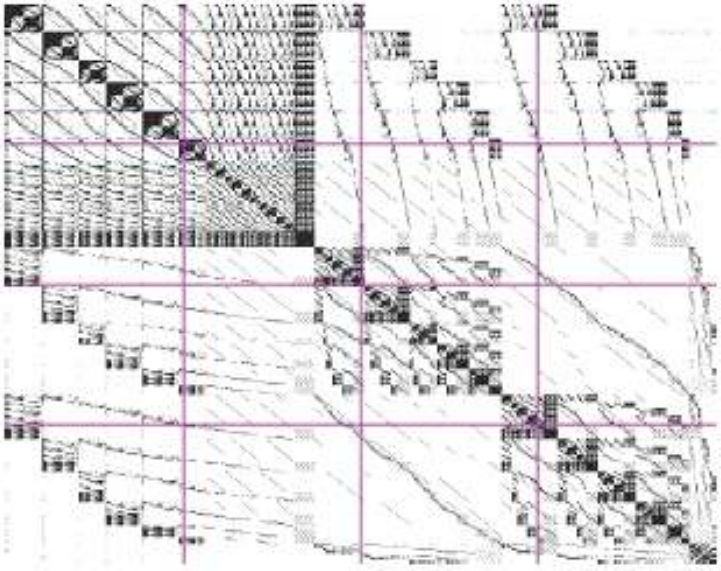

on single processor

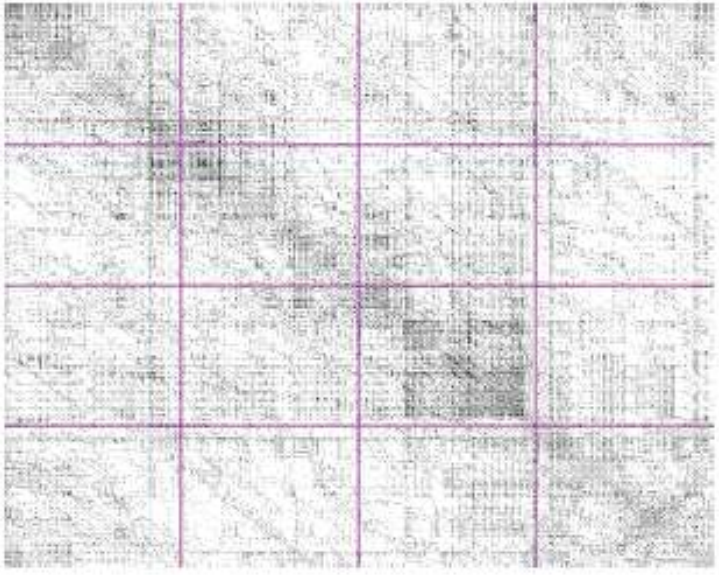

on 10 processors

- excellent load-balancing
- however, no (apparent) structure in sparse matrix

MFDn – Communication patterns

- Store lower half of symmetric matrix, distributed over $n = d \cdot (d+1)/2$ processors with d "diagonal" proc's
- Store Lanczos vectors on one of $(d+1)/2$ groups of d proc's
- Communication pattern matrix-vector multiplication

lower triangle

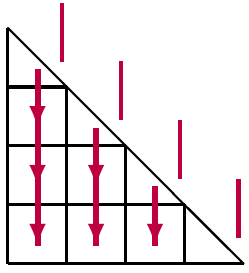

BCast(x)

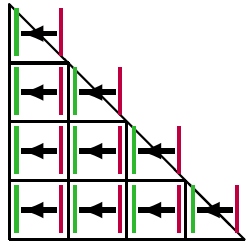

$y \leftarrow Ax$

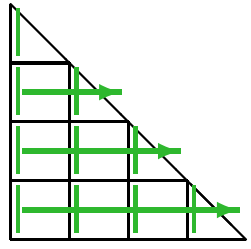

Reduce(y)

upper triangle

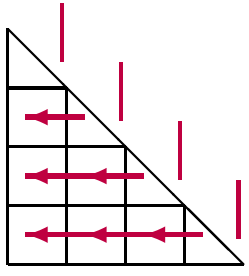

BCast(x)

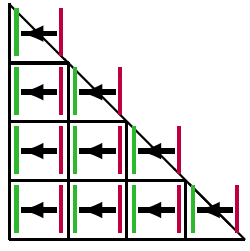

$y \leftarrow A^T x$

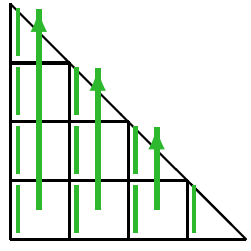

Reduce(y)

Speedup of MFDn on Cray XT4 (Franklin)

Maris, Sosonkina, Vary, Ng, and Yang, *Scaling of ab-initio nuclear physics calculations on multicore computer architectures*, Procedia Computer Science 1, May 2010, pages 97-106 (ICCS 2010)

- strong scaling, ^{12}C, $N_{\max} = 4$, 2-body interactions
- dimension $1.1 \cdot 10^6$, # nonzero m.e. $280 \cdot 10^6$
- memory for storing matrix: 2.3 GB
- speedup $T_{1\text{ CPU}}/T_{\#\text{ CPU's}}$

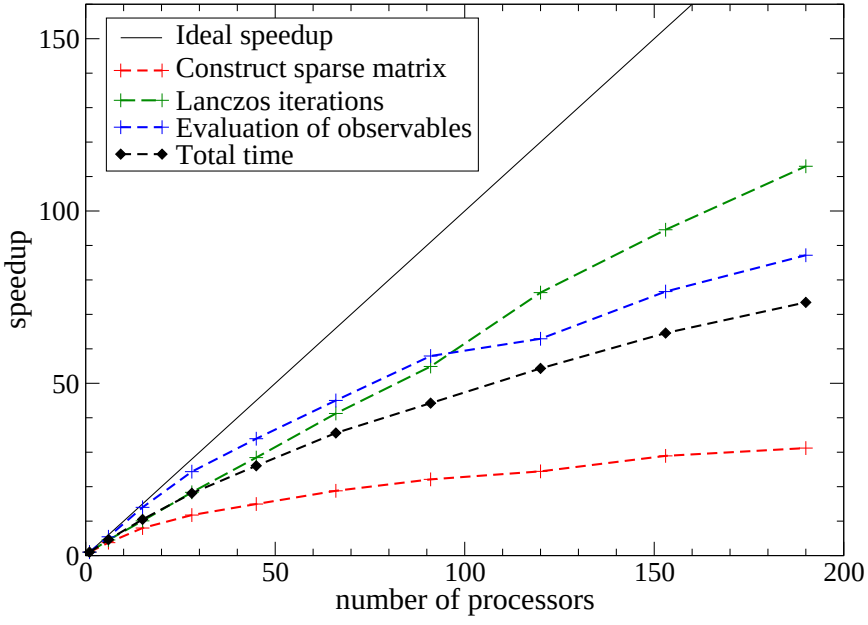

using MPI only

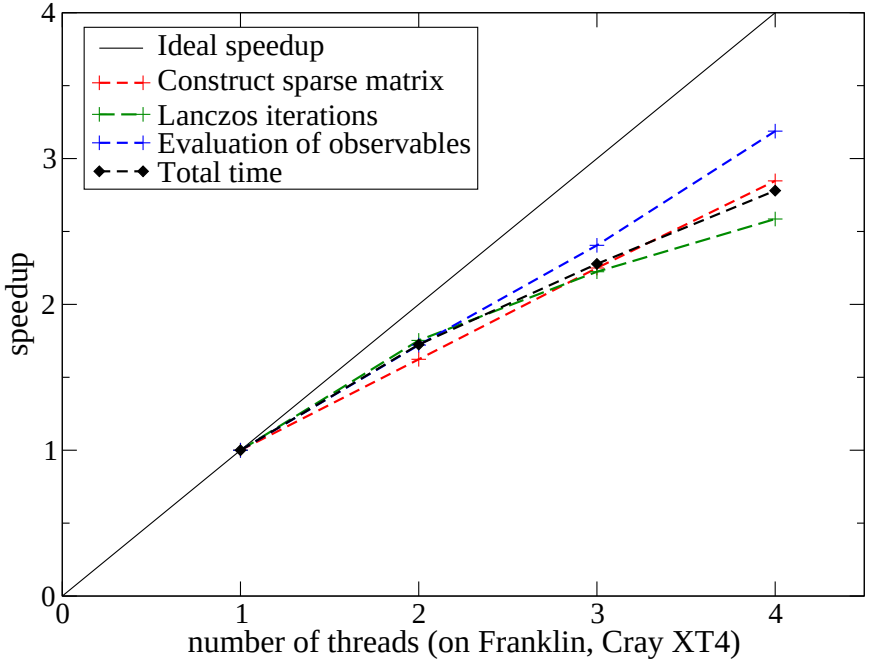

using 1 to 4 threads on single node

Total CPU time of MFDn on Cray XT4 with hybrid MPI/OpenMP

- For application scientist, time to completion, or CPU resource units used, is more important than speedup

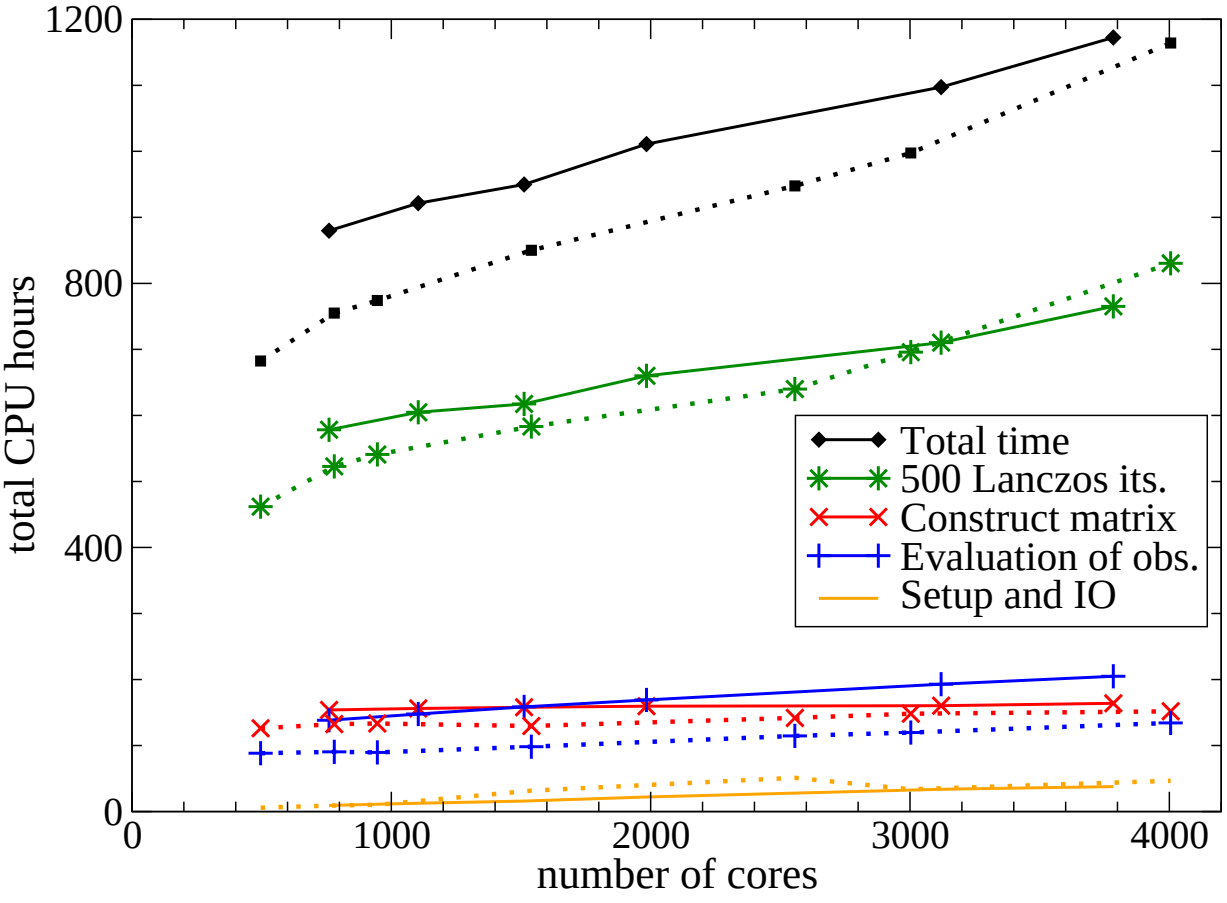

^{6}Li, $N_{\max} = 12$,
2-body interactions
on Franklin (NERSC)

solid: hybrid MPI/OMP
1 MPI PE per node
with 4 threads

dotted: pure MPI
4 MPI PE's per node

dimension $49 \cdot 10^6$
nonzero m.e. $74 \cdot 10^9$
memory for matrix: 600 GB

- Pure MPI more efficient than hybrid MPI/OpenMP for this case

Hybrid MPI/OpenMP more efficient as problem size grows

^{14}N, $N_{\max} = 8$, 2-body interactions, on Franklin (XT4) and Jaguar (XT5)

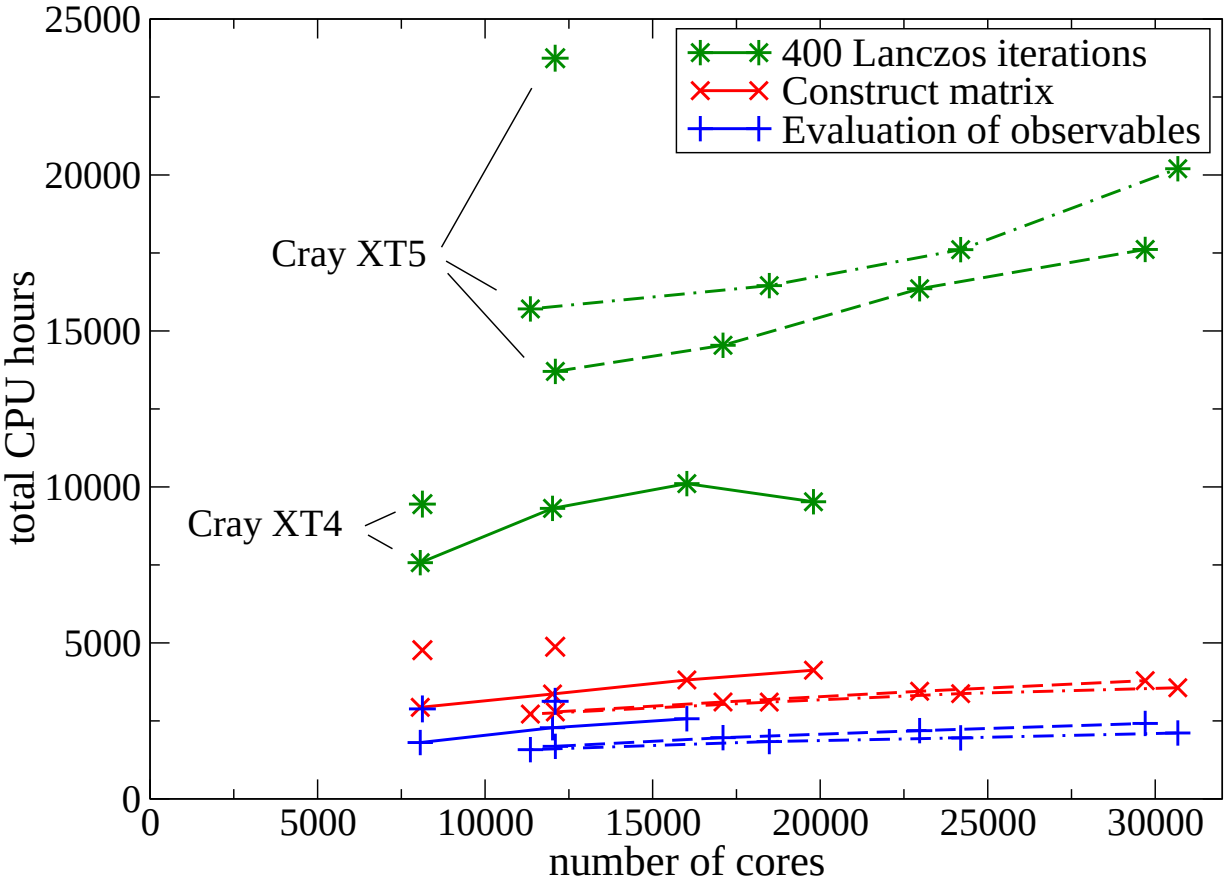

solid: hybrid OMP/MPI
1 MPI PE per node
with 4 threads (XT4)

dashed: hybrid OMP/MPI
1 MPI PE per NUMA node
with 6 threads (XT5)

dot-dashed: hybrid OMP/MPI
1 MPI PE per compute node
with 12 threads (XT5)

dimension $1.1 \cdot 10^9$
nonzero m.e. $1 \cdot 10^{12}$
memory for matrix: 8 TB

For comparison: symbols at 8,128 (XT4) and at 12,090 (XT5) cores
pure MPI with 1 MPI PE per core

Quad core vs. Hex core performance on Cray XT5 at ORNL

^{14}Be, dimension $2.8 \cdot 10^9$, # nonzero m.e. $2.8 \cdot 10^{12}$, memory for matrix: 24 TB

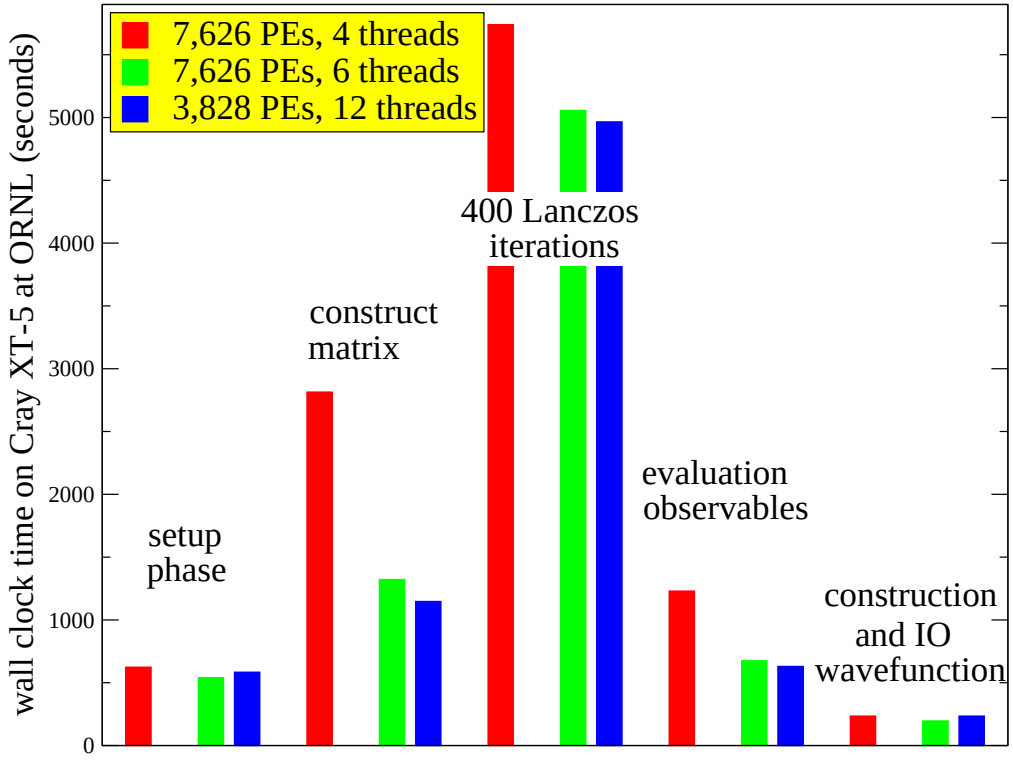

Sept/Oct 2009:
Upgrade Jaguar at ORNL
to Istanbul processor
expect improvement 1.7

quad-core $\rightarrow$ hexcore
clockspeed: 2.3 GHz $\rightarrow$ 2.6 GHz
L3 cache: 2 MB $\rightarrow$ 8 MB
increased memory bandwith
memory kept at 8 GB per proc
interconnect kept the same

- Construction of matrix time improved by factor of 2.13, probably due to increased cache and/or memory bandwith
- Lanczos iteration time improved by only factor of 1.14, most likely limited by MPI communication time

Interface between MFDn and NEWUOA/VTDIRECT

- optimize Wood–Saxon basis function parameters by minimizing calculated ground state energy for given interaction
- optimize (effective) NN interaction by minimizing RMS difference between set of calculated observables and experimental values (absolute binding energies and/or excitation energies)

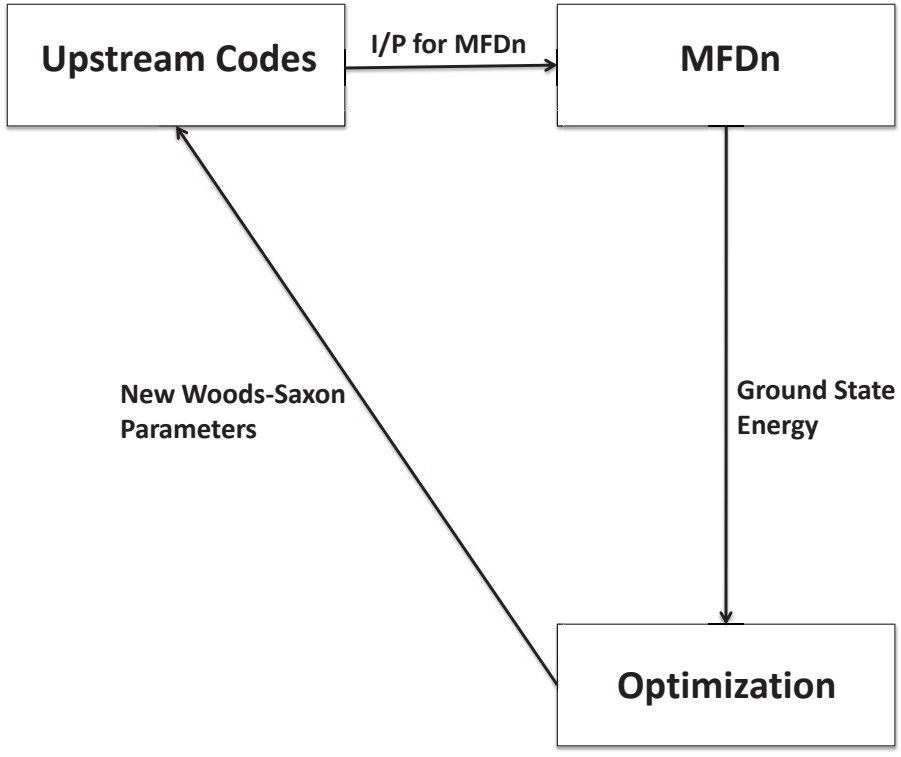

Checkpointing

- Execution of MFDn expensive
  need to be able to checkpoint/restart optimization of parameters
- Checkpointing implemented in interface MFDn and NEWUOA
- Checkpointing overhead is insignificant compared to the benefit of saved computation for expensive function evaluations

  T_{fe}: function evaluation; T_{sv}: save checkpiont data; T_r: recovery time

Nucleus, N_{max}, N_{shell}	T_{fe} (seconds)	T_{sv} + T_r (μ seconds)
^{4}He, 6, 14	837	778
^{4}He, 8, 14	880	777
^{4}He, 10, 14	1100	800

note difference in scale between columns

- Note: currently no checpoint capability with MFDn Version 13
  but Version 14 has checkpoint/restart capabilities

Results with Wood–Saxon basis – 4He

- Does a more realistic basis exhibit faster convergence, in particular for RMS radii?
  E.g. Wood–Saxon basis

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

- Test with NN potential JISP16
  - optimize WS parameters by minimizing ground state energy
  - slight improvement of variational bound ground state energy
  - no improvement in convergence RMS radius

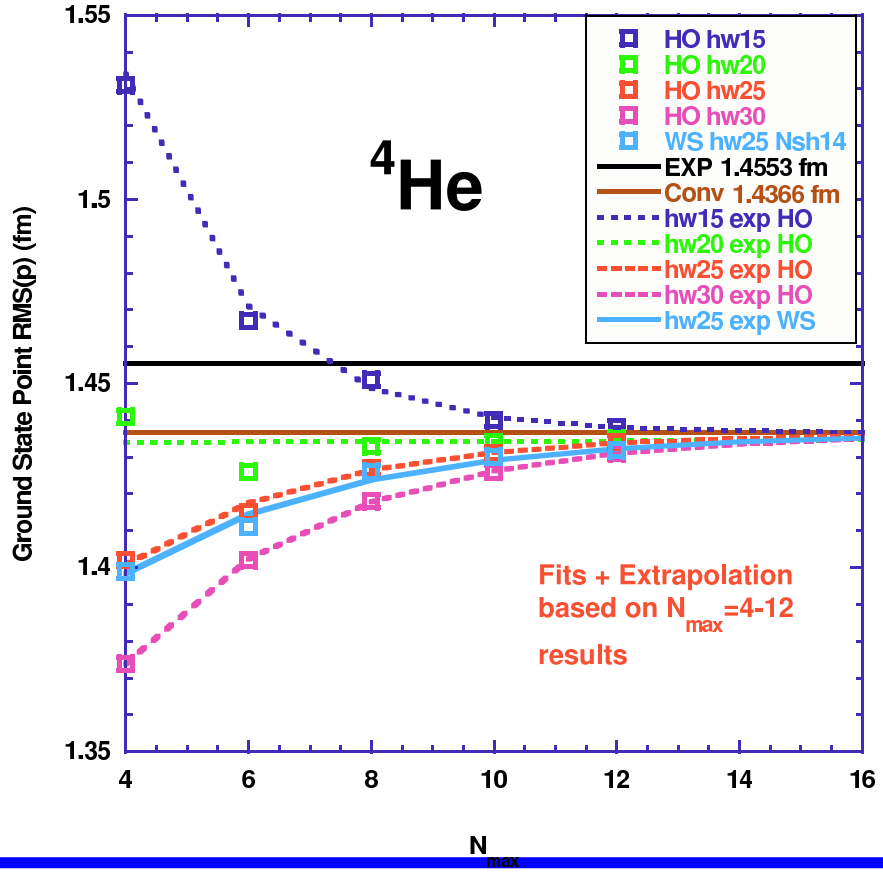

Results with Wood–Saxon basis – 6He

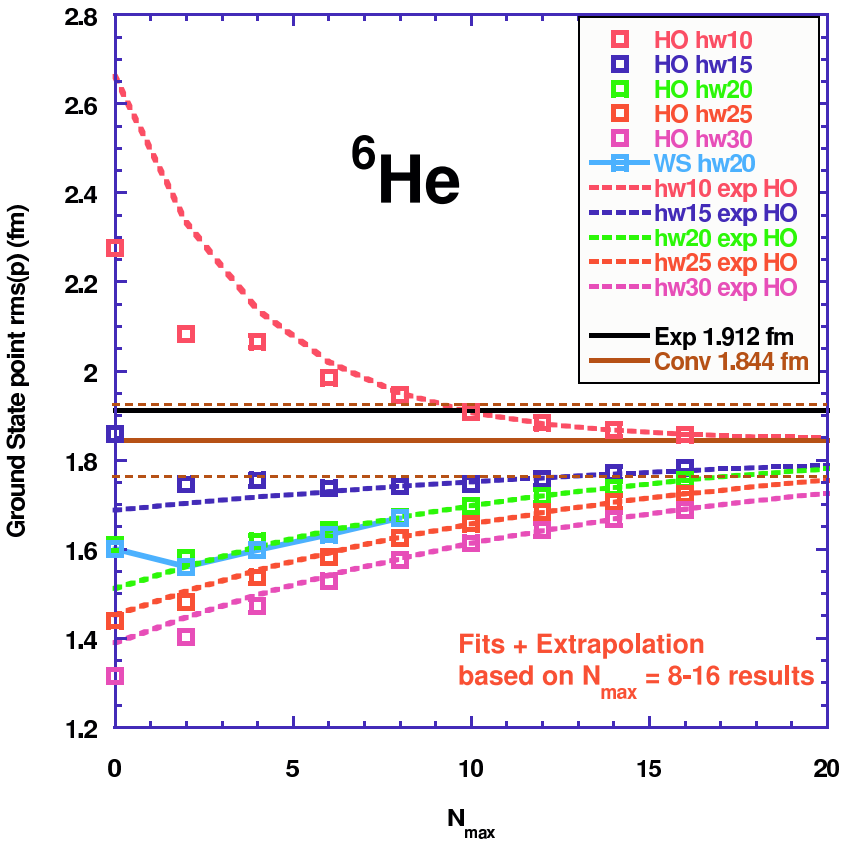

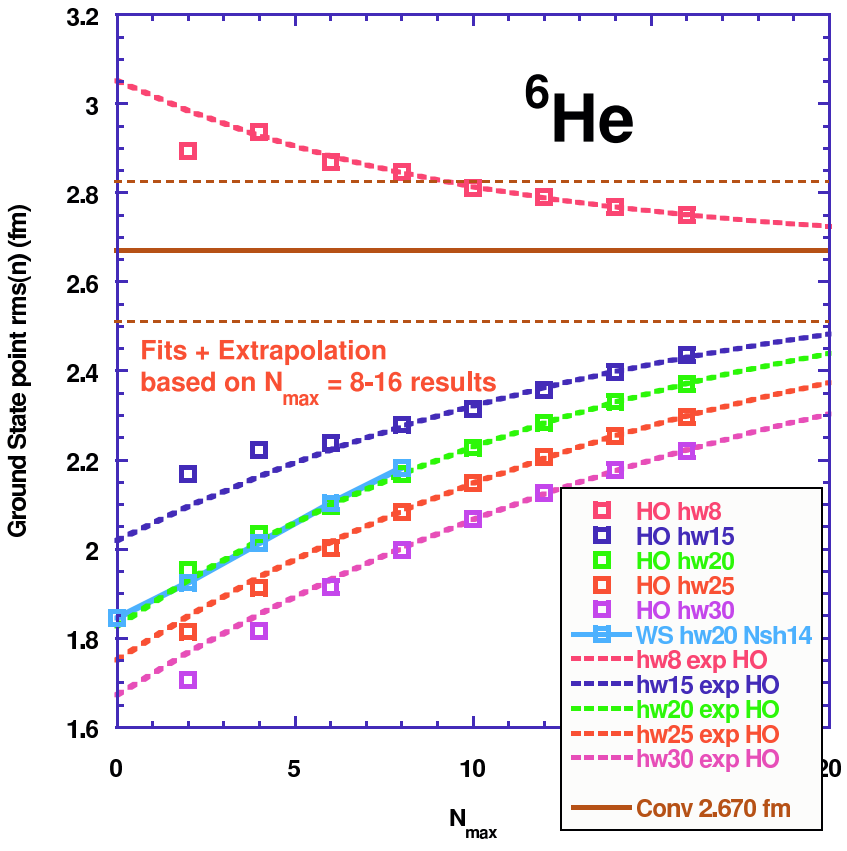

- no improvement in convergence RMS radius of ^{6}He
  optimal WS basis similar to optimal HO basis
- both proton radius and neutron radii extrapolated using exponential
- for more details, see PhD thesis Negoita (defense next week)

Results with Wood–Saxon basis – 12C

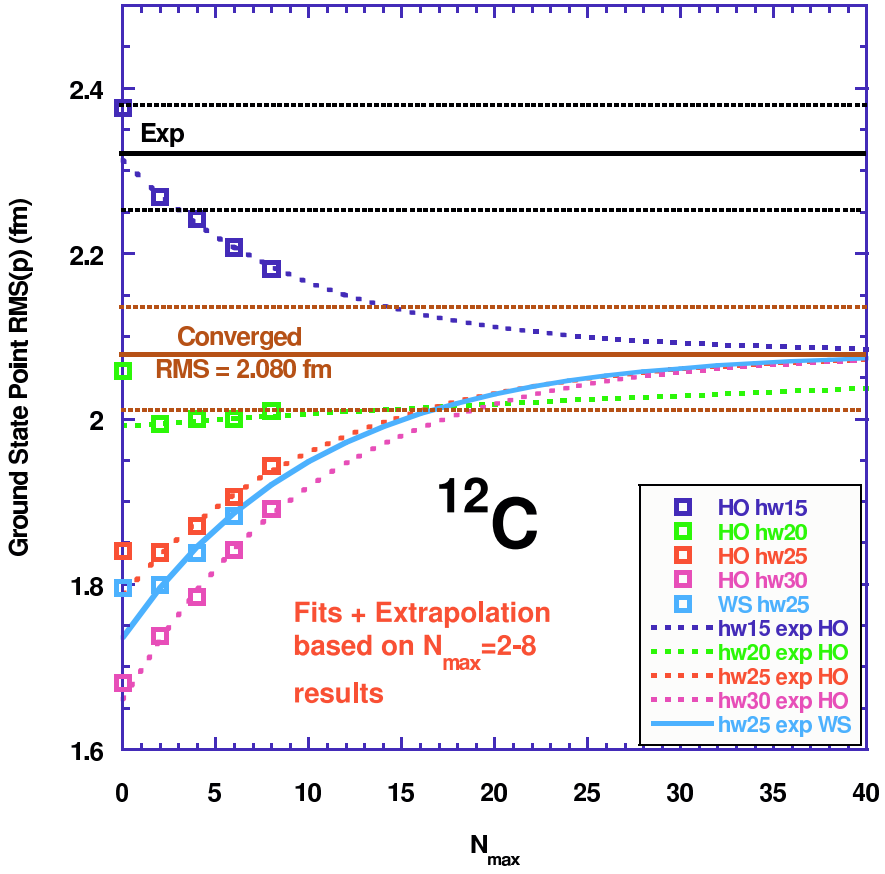

- no improvement in convergence RMS radius of ^{12}C
  optimal WS basis similar to optimal HO basis
- calculated radius significantly below experimental radius
- for more details, see PhD thesis Negoita (defense next week)

NEW One-Body Density Matrix Elements from MFDn

- Static OBDMEs (for protons and neutrons seperately)
- Transition OBDMEs to/from one reference state
  (not necessarily the ground state)
- Current I/O format
  - one file with quantum numbers
    i_α(S.P. orbital), i_β(S.P. state), n_α, l_α, $2j_\alpha$, n_β, l_β, $2j_\beta$, $2mj$, $2tz$
    for all potentially nonzero OBDME's (i.e. allowed by symmetries)
  - one file for each static and transition OBDME,
    with all all potentially nonzero OBDME's
  - note: NOT the reduced OBDME's, because we plan to enable deformed (axially symmetric) external fields (soon)
- Used for calculation of one-body observables within MFDn
- Available for collaborators for postprocessing
  - storage/archiving of OBDMEs – see prototype LCCI database
- Question: useful for visualization?

One-Body Density Matrix for ^{8}He

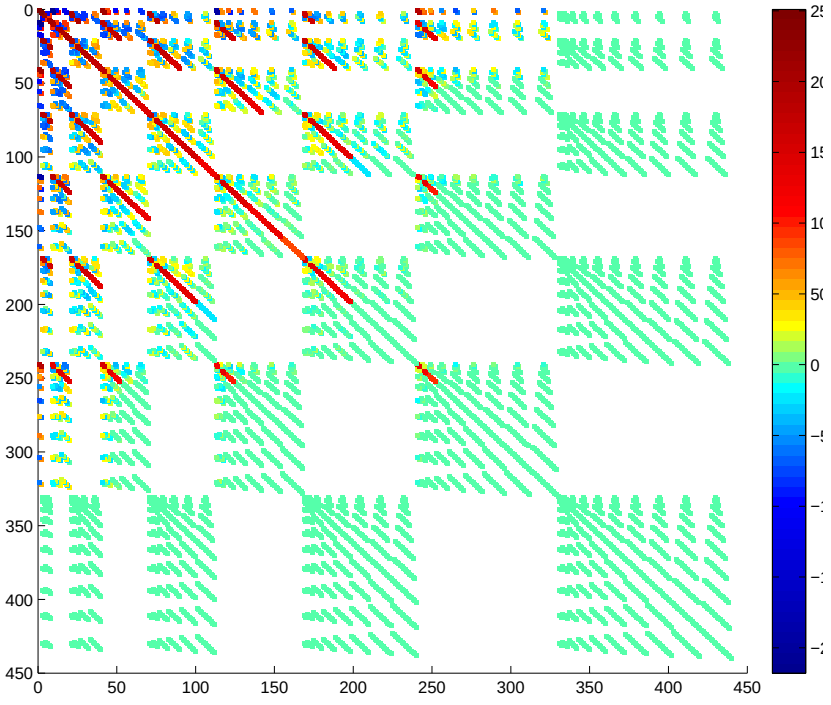

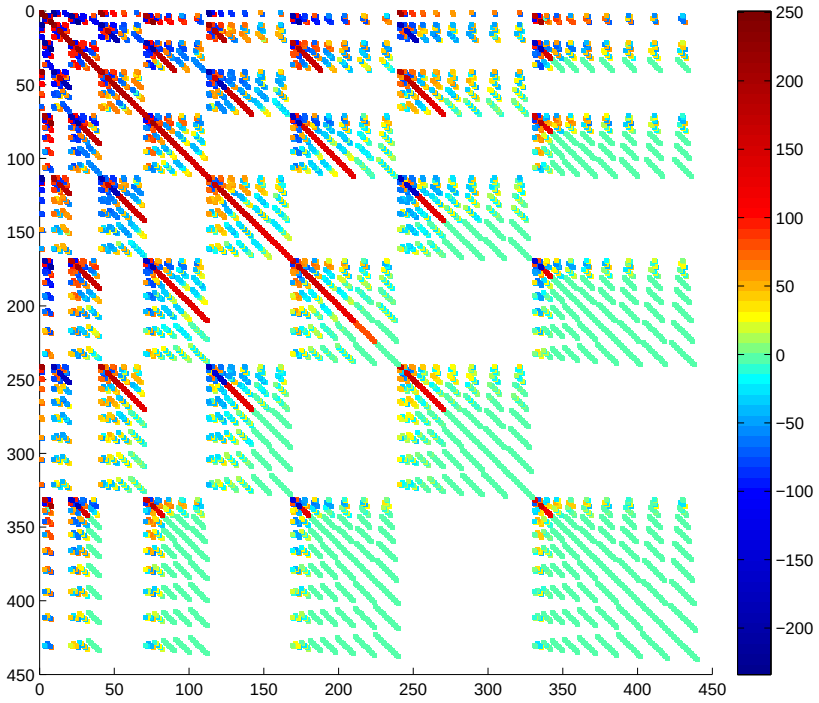

- attempt at visualization of OBDME for ^{8}He
- $N_{\max} = 8$, so 10 S.P. shells

One-body densities for 12C with JISP16

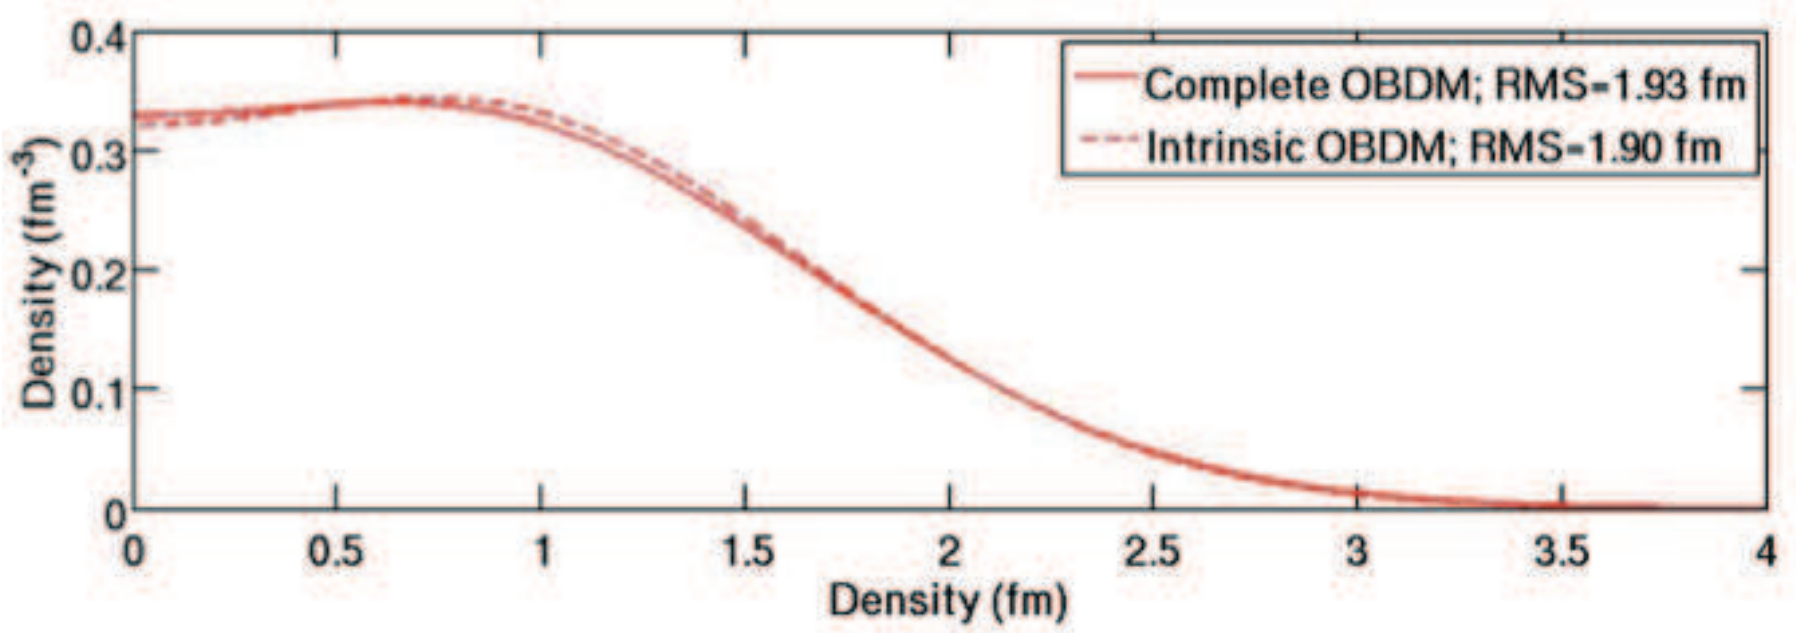

- MFDn use single-particle coordinates, not relative coordinates
- Can de-convolute CoM and intrinsic coordinates for matter density but not proton and neutron density seperately
- Difference intrinsic density and complete density small
- With TRDENS intrinsic OBDMEs available (using 2-body operator)

One-body densities for 12C with JISP16

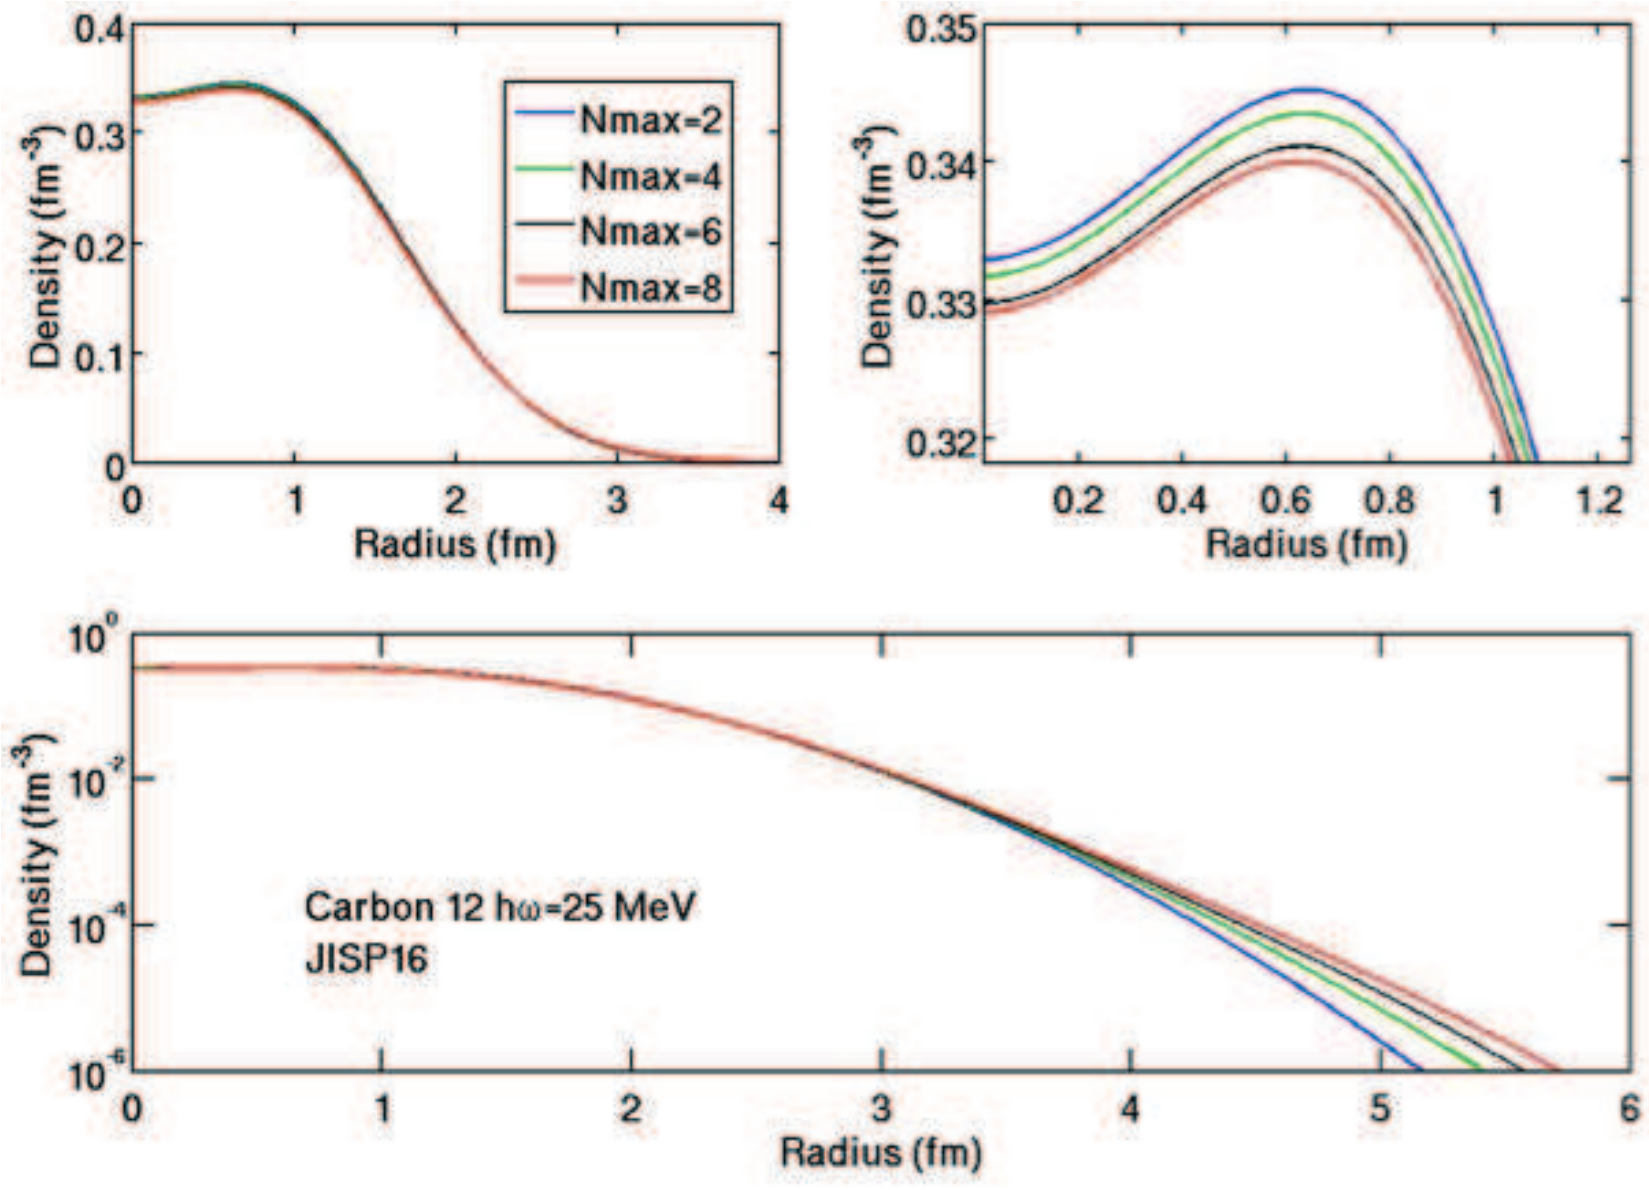

One-body densities – projects in progress

- Evaluation of static and transition one-body density matrices and electroweak amplitudes from the SM and, together, evaluate the ^{12}C$(\nu,\nu')^{12}$C cross section needed for long-baseline neutrino mixing experiments
  (w. Harry Lee, ANL, in progress)

- Ready for production of one-body density matrices of low-lying states of p-shell nuclei for comparison with, and validation of, DME/DFT methods
  - Which nuclei, which interactions?
    - NN potentials: JISP16, Minnesota, N3LO (renormalized)
    - NNN potentials: N2LO (either SRG or LSO renormalized)
  - With external fields for protons and/or neutrons?

- Evaluation of ground state one-body density matrix of neutron droplets in HO external field using Minnesota NN potential for comparison with, and validation of, DME methods
  (w. OSU/MSU/ORNL, in progress)

Neutron droplets

Preliminary results – DME plot from Markus Kortelainen
Density for 8 neutrons in a HO trap with Minnesota NN interaction

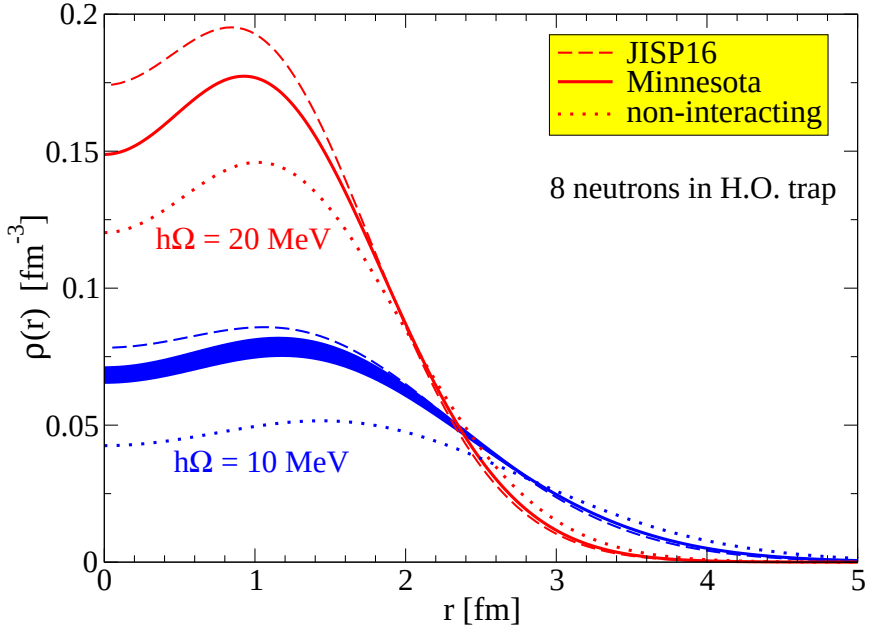

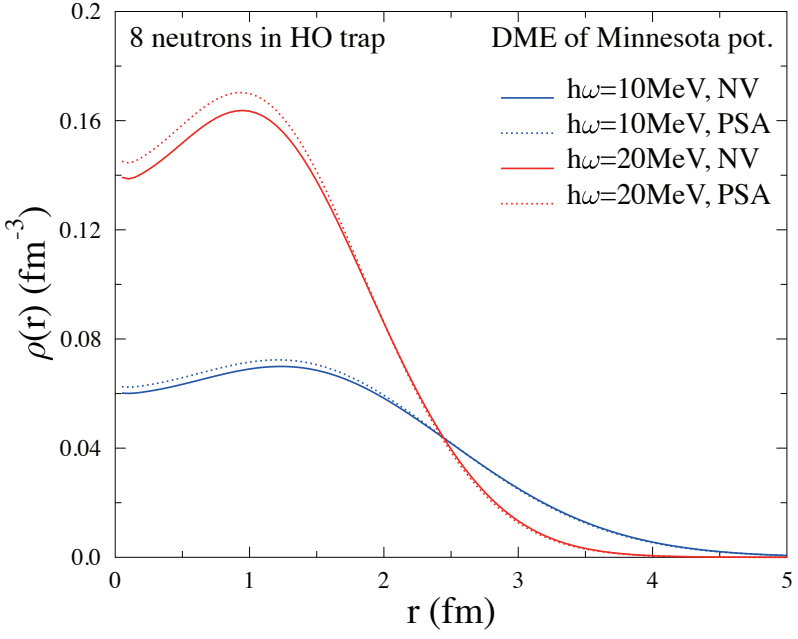

- qualitative agreement between NCFC and DME calculations
- more details discussed this afternoon

Neutron droplets

Preliminary results – DME plot from Markus Kortelainen
Density for 8 neutrons in a HO trap with Minnesota NN interaction

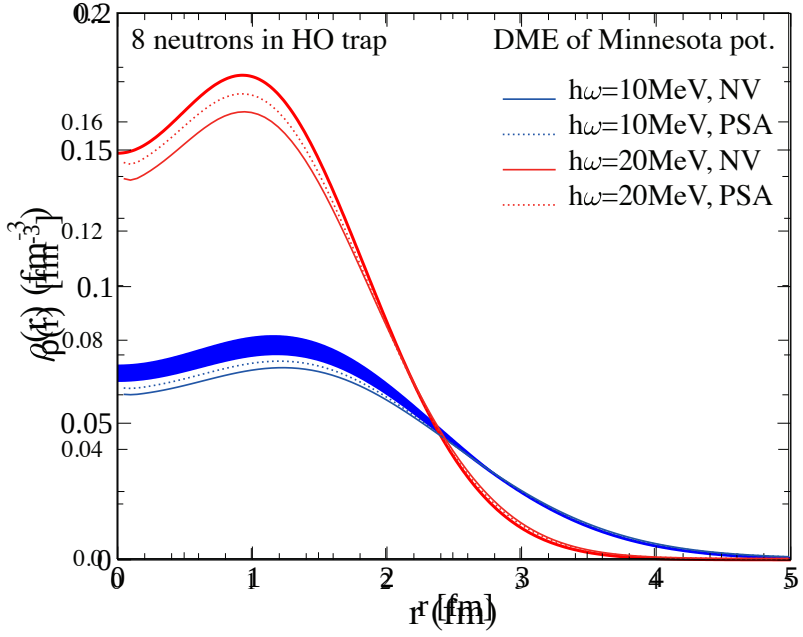

- qualitative agreement between NCFC and DME calculations
- more details discussed this afternoon

Comparison between methods and interactions – 8 neutrons

interaction	method	field (MeV)	E_{total} (MeV)	E_{internal} (MeV)	radius (fm)
none		10.0	180.00	90.00	3.0554
Skyrme	EDF	10.0	110.93		
AV8' + UIX	GFMC	10.0	135.76(4)	64.5(1)	2.72
AV8' + UIX	AFDMC	10.0	134.75(9)		
JISP16	NCFC	10.0	132.5(1)	63.4(4)	2.675(15)
Minnesota	NCFC	10.0	138.(1)	66.(1)	2.73(3)
Minnesota	DME-NV	10.0	144.50	69.48	2.79
Minnesota	DME-PSA	10.0	143.38	69.57	2.77
none		20.0	360.00	180.00	2.1605
JISP16	NCFC	20.0	278.3(5)	126.(1)	1.99(2)
Minnesota	NCFC	20.0	290.7(2)	131.3(3)	2.033(4)
Minnesota	DME-NV	20.0	304.18	136.21	2.09
Minnesota	DME-PSA	20.0	299.91	135.53	2.06

Preliminary results

Many Fermion Dynamics (MFD) nuclear physics – Version 14

Completely new version of MFDn *under development*

- Simplified input (compared to Version 13)
- Different storage and distribution of many-body basis states
  - round robin distribution of groups of states
    inspired by Total-J code
  - advantage: retain some of block structure in sparse matrix
    quickly determine blocks of potentially nonzero matrix elements
  - challenge: load-balancing
- Capable of check-pointing: each Lanczos vectors written to file
- Capable of "partial re-compute on the fly", reducing memory needs
  - $N_{\max} = 10$ runs with 2-body forces for ^{12}C and ^{14}C
  - $N_{\max} = 8$ runs with 3-body forces throughout the p-shell
- Has run succesfully on 220,000 cores on JaquarPF
- Same 1- and 2-body observables as Version 13 *under development*

Ab initio deliverables – Year 4

- ^{14}C to ^{14}N β decay
  - several large runs with 3-body forces performed at ORNL
  - draft paper – in progress
- OBDME's for nuclei – in progress
  - investigation of ^{12}C
    - intrinsic density, E2 transitions, Hoyle state
    - neutrino cross sections (w. Harry Lee – ANL)
  - use of ab initio OBDME's to inform DFT's ?
- Neutron droplets in external fields – in progress
  - new results with JISP16, Minnesota
  - validation of DME calculations w. Minnesota potential
- Use of realistic basis states
  - Wood–Saxon basis small improvement in convergence of E_{gs}
- Ab initio reactions using J-matrix methods
  - progress thanks to parallel Total-J code

Roadmap – Remainder of Year 4 and Plans for Year 5

- Code development MFDn Version 13
  - axially deformed HO & WS external fields (yr 4, ISU)
  - script, database, interfaces w. up- and down-stream codes, work-flow management system? (e.g. Kepler) (yr 4-5, LCCI)
- Code development MFDn Version 14 (yr 4-5, ISU/LBNL/ORNL)
  - improve single-processor performance and scalability
  - flexible switch between “in-core” and “on-the-fly”
- Set of neutron properties for DFT/DME communities
  - Minnesota potential, 4 to 20 neutrons (yr 4, ISU)
  - SRG-evolved chiral 2- and 3-body forces (yr 5, ISU/LLNL)
  - deformed HO/WS external fields, as needed/requested (yr 5)
- Chiral 2- and 3-body runs for A = 9 through 16 (yr 4-5, ISU/LLNL)
- Continue investigation of realistic basis functions (ISU/AL)
- Investigate alternative eigensolver, block algorithms, improvements Total-J code (yr 5, LBNL/ISU)