Recent NCFC results for light nuclei with SRG evolved chiral interactions

Pieter Maris
pmaris@iastate.edu
Iowa State University

SciDAC project – UNEDF
PIs: Rusty Lusk (ANL), Witek Nazarewicz (ORNL/UT)
http://www.unedf.org

PetaApps award
PIs: Jerry Draayer (LSU), Umit Catalyurek (OSU)
Masha Sosonkina, James Vary (ISU)

INCITE award – Computational Nuclear Structure
PI: James Vary (ISU)

NERSC CPU time

Configuration Interaction Methods in a nutshell

- Expand wave function in basis states $|\Psi\rangle = \sum a_i |\psi_i\rangle$
- Express Hamiltonian in basis $\langle\psi_j|\hat{\mathbf{H}}|\psi_i\rangle = H_{ij}$

$$\hat{\mathbf{H}} = \hat{\mathbf{T}}_{\text{rel}} + \Lambda_{CM}\left(\hat{\mathbf{H}}_{CM}^{H.O.} - \frac{3}{2}\hbar\omega\right) + \sum_{i<j} V_{ij} + \sum_{i<j<k} V_{ijk} + \ldots$$

- Pick your favorite potential
  - Argonne potentials: AV8, AV18 (plus Illinois NNN interactions)
  - Bonn potentials
  - Chiral NN plus NNN interactions
  - ...
  - JISP16 (phenomenological NN potential)
  - ...

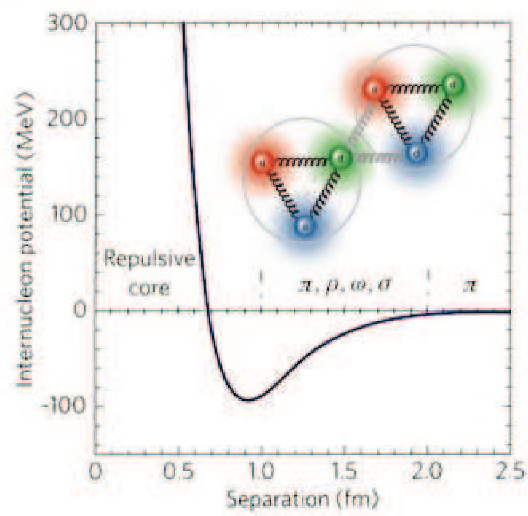

- Optional: Renormalize V in order to improve convergence

CI calculation – sparse matrix problem

- Expand wave function in basis states $|\Psi\rangle = \sum a_i |\psi_i\rangle$
- Express Hamiltonian in basis $\langle\psi_j|\hat{\mathbf{H}}|\psi_i\rangle = H_{ij}$
- Diagonalize sparse real symmetric matrix H_{ij}
- Complete basis $\longrightarrow$ exact result
  - caveat: complete basis is infinite dimensional
- In practice
  - truncate basis
  - study behavior of observables as function of truncation
  - infinite basis space limit: No-Core Full Configuration (NCFC)
- Computational challenge
  - construct large ($10^{10} \times 10^{10}$) sparse symmetric real matrix H_{ij}
  - obtain lowest eigenvalues & eigenvectors w. Lanczos algorithm

Intermezzo: Extrapolation Techniques

Challenge: achieve numerical convergence for no-core Full Configuation calculations using finite model space calculations

- Perform a series of calculations with increasing N_{max} truncation, while keeping everything else fixed!
- Extrapolate to infinite model space $\longrightarrow$ exact results
  - binding energy: exponential in N_{max}

$$E^N_{\text{binding}} = E^{\infty}_{\text{binding}} + a_1 \exp(-a_2 N_{\text{max}})$$

  - use 3 or 4 consecutive N_{max} values to determine $E^{\infty}_{\text{binding}}$
  - use $\hbar\omega$ and N_{max} dependence to estimate numerical error bars
  - need at least $N_{\text{max}} = 8$ for meaningfull extrapolations

PM, Vary, Shirokov, arXiv:0808.3420 [nucl-th], PRC79, 014308 (2009)

Intermezzo: Extrapolation Techniques

Challenge: achieve numerical convergence for no-core Full Configuation calculations using finite model space calculations

- Perform a series of calculations with increasing N_{max} truncation, while keeping everything else fixed!
- Extrapolate to infinite model space $\longrightarrow$ exact results

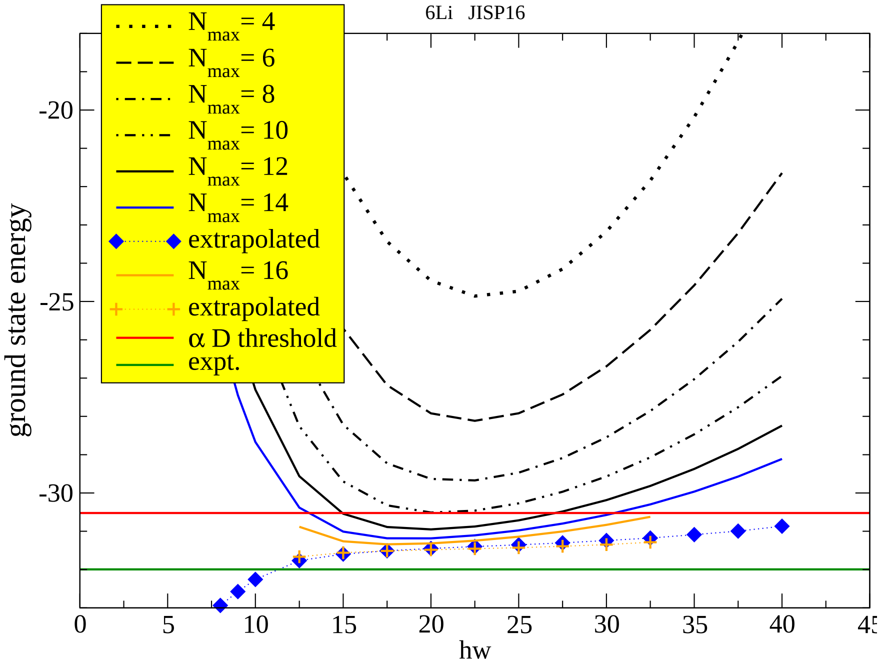

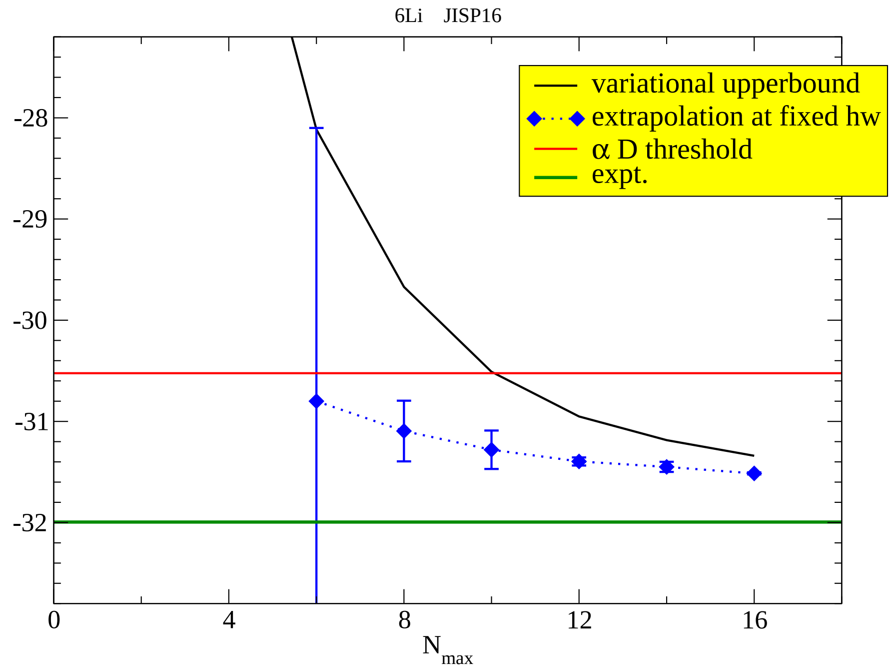

PM, Vary, Shirokov, arXiv:0808.3420 [nucl-th], PRC79, 014308 (2009)

CI calculations – main challenges

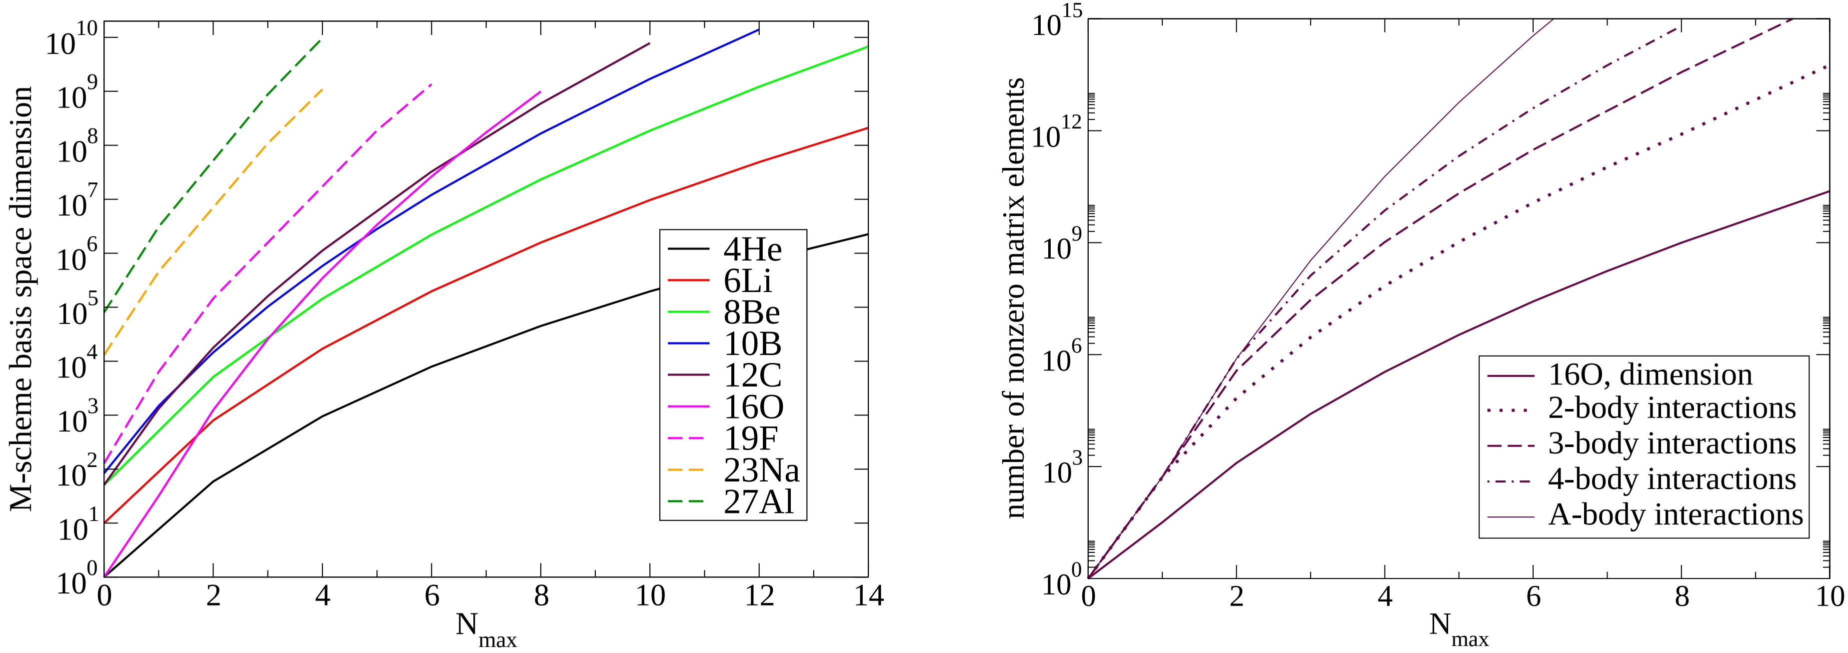

- Single most important computational issue:
  exponential increase of dimensionality with increasing $N_{\max}$
- Additional computational issue:
  sparseness of matrix / number of nonzero matrix elements
- Extrapolation to infinite basis requires $N_{\max} \geq 8$

Many Fermion Dynamics – nuclear physics

- Platform-independent, hybrid OpenMP/MPI, Fortran 90
- Can in principle handle arbitrary N-body interactions
  however input format only specified for 2- and 3-body interactions
- Generate many-body basis space
  subject to user-defined truncation and symmetry constraints
- Construction of many-body matrix H_{ij}
  - determine which matrix elements can be nonzero
    based on quantum numbers of underlying single-particle states
  - evaluate and store nonzero matrix elements
    in compressed row/column format
- Obtain lowest eigenpairs using Lanczos algorithm
  - typical use: 10 to 20 lowest eigenvalues and eigenvectors
  - typically need $\sim$ 500 to $\sim$ 1000 Lanczos iterations
- Calculate select one- and two-body observables
- One-body density matrices and wavefunctions available
  as input scattering and reaction calculations

MFDn – 2-dimensional distribution of matrix

- Real symmetric matrix: store only lower (or upper) triangle distributed over $n = d \cdot (d+1)/2$ processors with d "diagonal" proc's
- In principle, we can deal with arbitrary large dimensions even if we cannot store an entire vector on a single processor
- Communication pattern matrix-vector multiplication

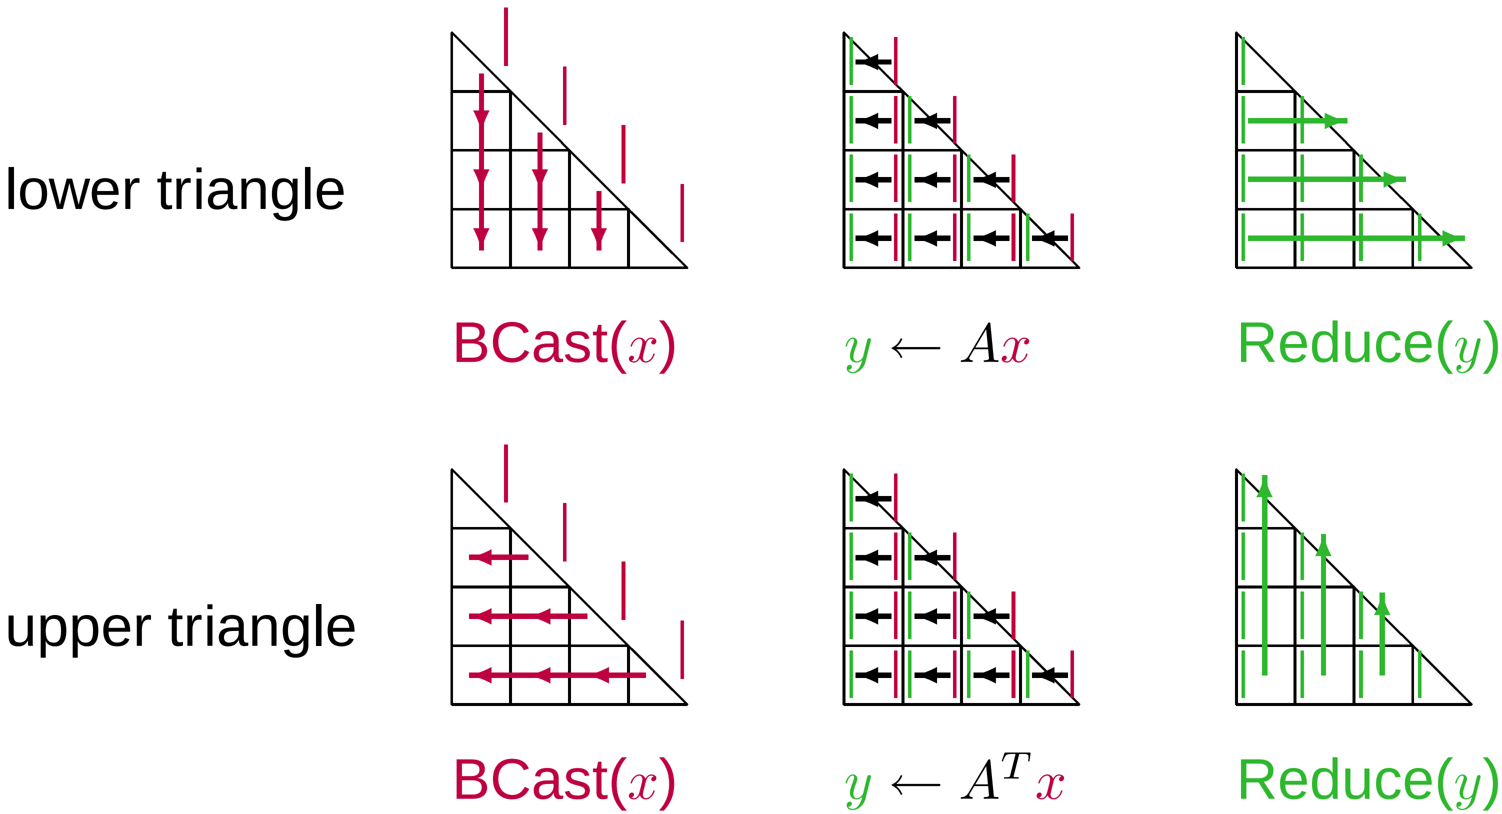

MFDn – communication patterns and processor layout

0				
5	1			
9	6	2		
12	10	7	3	
14	13	11	8	4

original layout

0			11	13
1	3			14
2	4	6		
	5	7	9	
		8	10	12

significantly improved performance

(Aktulga 2012)

- SpMV
  - Bcast along columns
  - Bcast along rows
  - local SpMV
  - local transpose SpMV
  - Reduce along columns
  - Reduce along rows
- Orthogonalization
  - Scatterv from diags to all procs
  - dot-products distributed over all procs

    Note: vectors and matrix in single precision,

    but accumulate dot-products in double precision
  - Gatherv from all procs to diags

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14

Strong Scaling of MFDn on Cray XT4

MPI strong scaling, ^{10}B, $N_{\max} = 6$, 2-body interactions

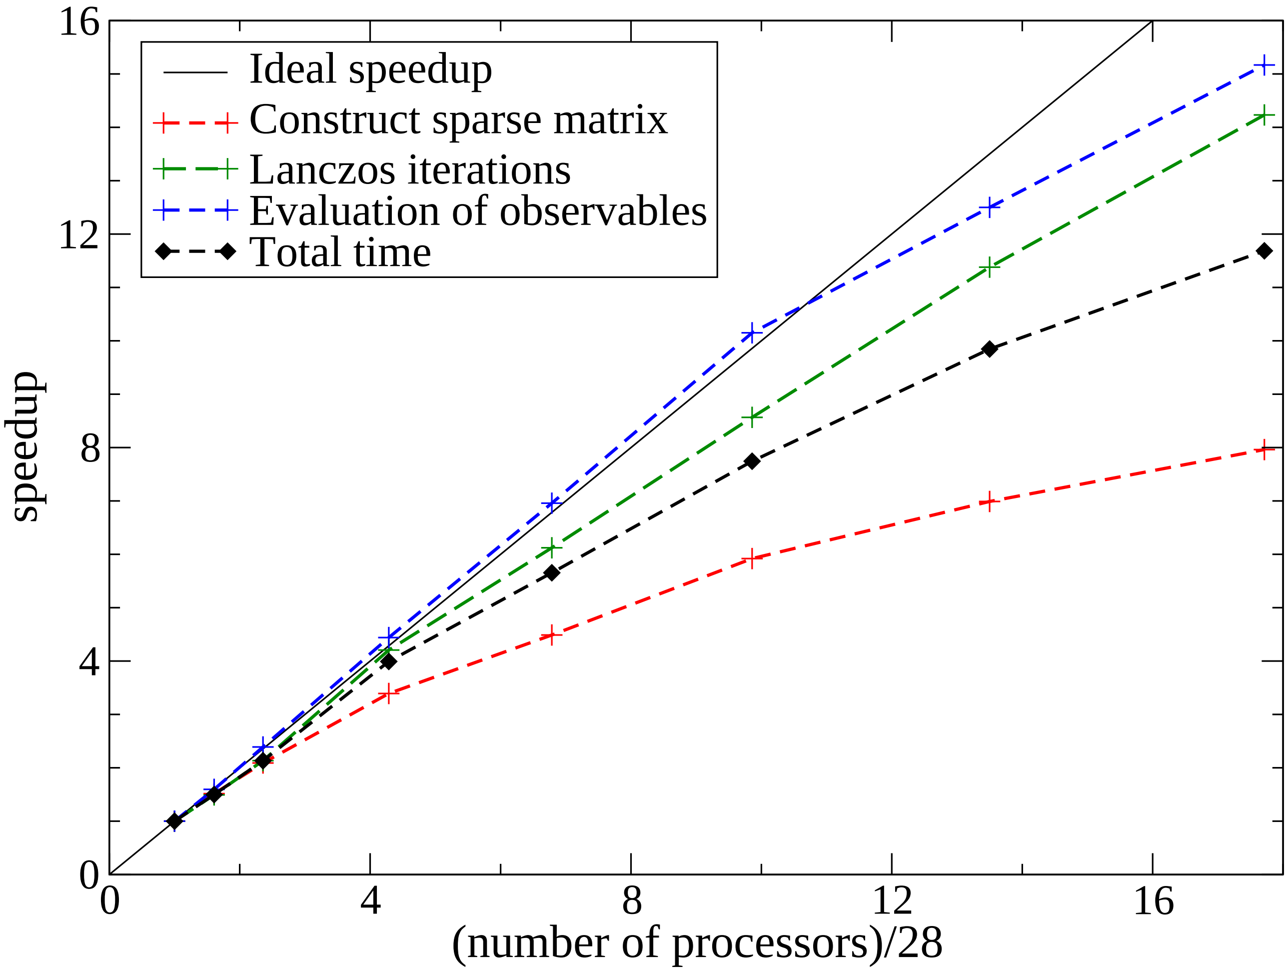

using 4 MPI procs per node on Franklin (NERSC)

$$\text{speedup} = \frac{T_{\text{28 PE's}}}{T_{\text{\# PE's}}}$$

dimension $10 \cdot 10^6$

\# nonzero m.e. $4.4 \cdot 10^9$

memory for storing matrix: 36 GB

Hybrid MPI and OpenMP

- Modern supercomputers: multi-core architectures
- Amount of communication and corresponding buffers becomes more demanding with increasing number of MPI procs
- Amount of common data, needed on all processors, increases with model space, in particular for 3-body interactions

$\Longrightarrow$ Need hybrid OpenMP / MPI version of MFDn

- Construction of matrix
  - columns distributed over threads using OMP directives
- Lanczos iterations
  - SpMV: columns distributed over threads using OMP directives
  - special treatment of transpose SpMV necessary in order to avoid race conditions
  - orthogonalization distributed using OMP directives
- Evaluation of observables distributed using OMP directives

Intermezzo: Memory requirements

- Consider matrix dimension D with NNZ nonzero matrix elements
- Consider hybrid MPI/OMP run on total of N_c cores,
  with N_p MPI processors and n_t OpenMP threads
  and N_d 'diagonal' processors $N_p = N_d\,(N_d+1)/2$
- Perfect load-balancing
  - local dimension $D/\sqrt{2N_p}$
  - local number of nonzero mat.els. NNZ/N_p
- Global memory footprint during Lanczos iterations
  - $4\,D\,N_{\text{iterations}}$ Lanczos vectors
  - $(8\,NNZ + 4\,D\sqrt{N_c/(2\,n_t)})$ sparse matrix in CCF
  - $4\,D\,\sqrt{N_c/2}\;(4/\sqrt{n_t} + \sqrt{n_t})$ local work arrays
- Multi-threading with 4 threads
  minimizes global memory footprint during Lanczos iterations

Strong Scaling of MFDn on Cray XK6

Hybrid MPI/OpenMP scaling for ^{10}B, $N_{\max} = 8$ with 3-body interactions

- Current bottleneck: M-scheme input interaction file of 33 GB
- Can be removed: coupled JT-scheme

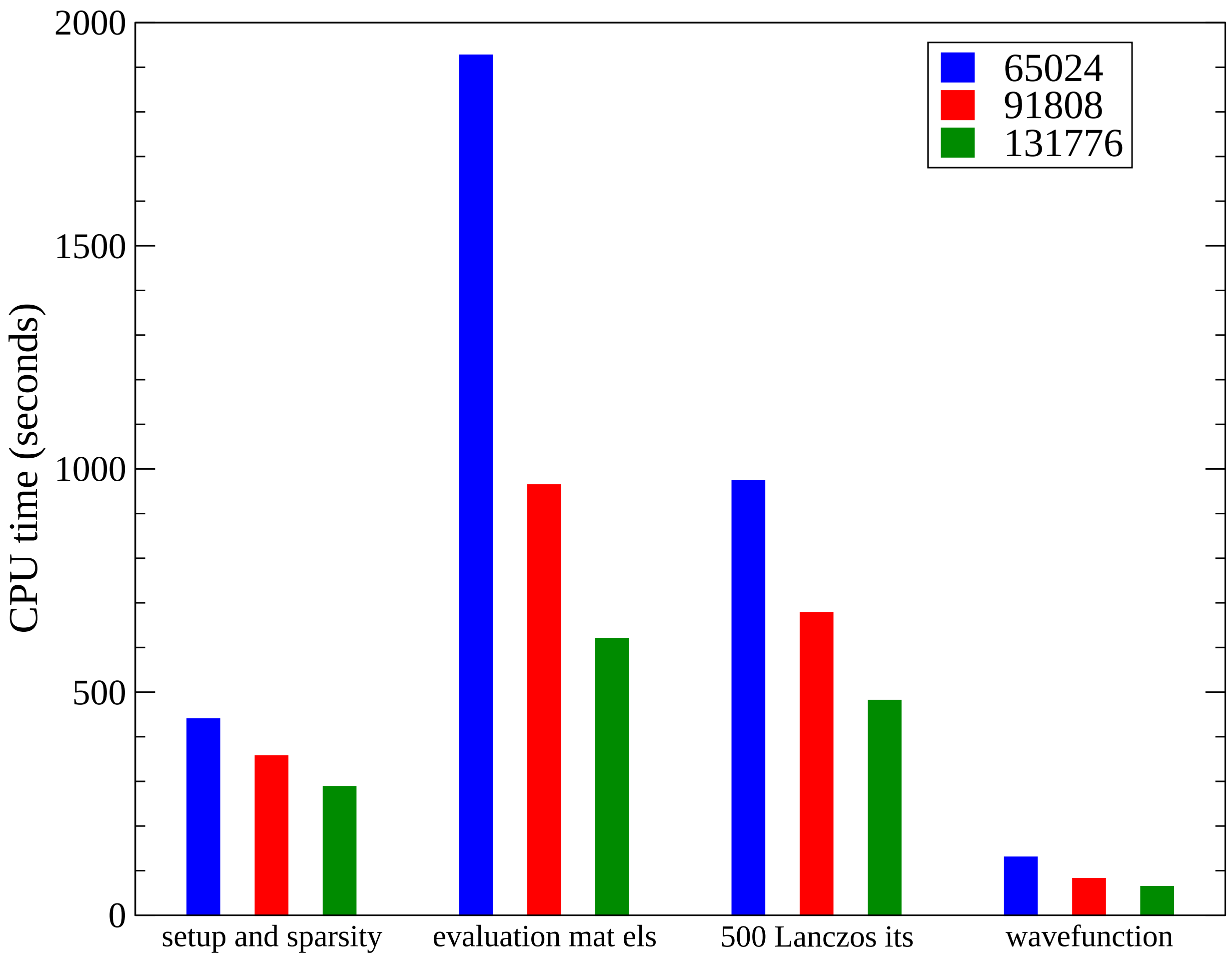

using 2 MPI procs per node
and 8 threads per MPI proc
on Jaguar (ORNL)

dimension $166 \cdot 10^6$
nonzero m.e. $5.43 \cdot 10^{12}$
memory for matrix: 44 TB

Benchmark calculations for light nuclei

Chiral NN interaction, SRG evolved to $\lambda = 1.5$, $\lambda = 2.0$, and $\lambda = 2.5$

- UNEDF benchmark project
  - initiated June 2007
  - CCSD (Hagen/Papenbrock) vs. NCFC (Maris/Vary)
  - unpublished results (June 2008)

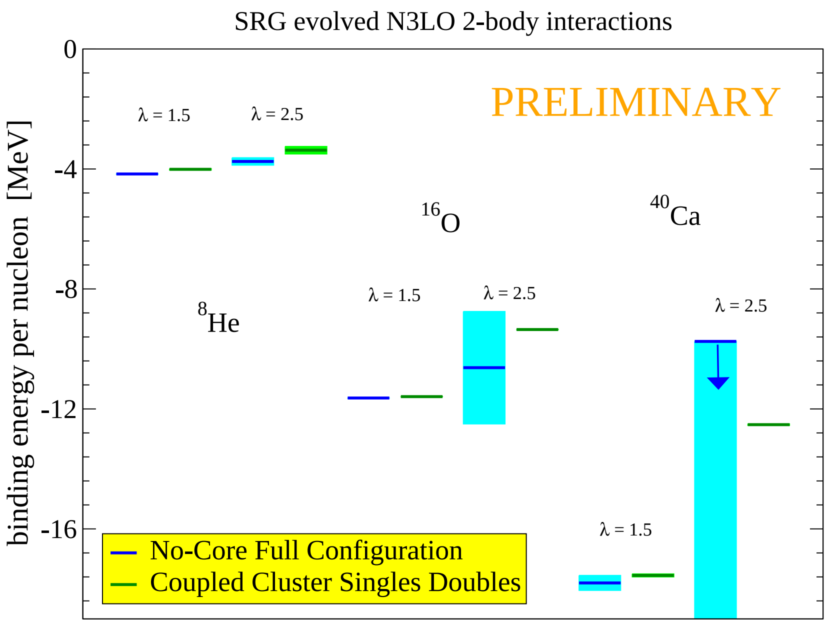

- Current benchmark project
  - Chiral NN interaction without Coulomb
  - Results plus estimate of all numerical uncertainties
  - CCSD(T), NCFC, it-NCSM, . . .
  - Data delivered: Nov 18, 2011 for gs energies and RMS radii of ^{8}He, ^{12}C, ^{16}O, ^{18}O, and ^{40}Ca

Ground state energy of ^{8}He

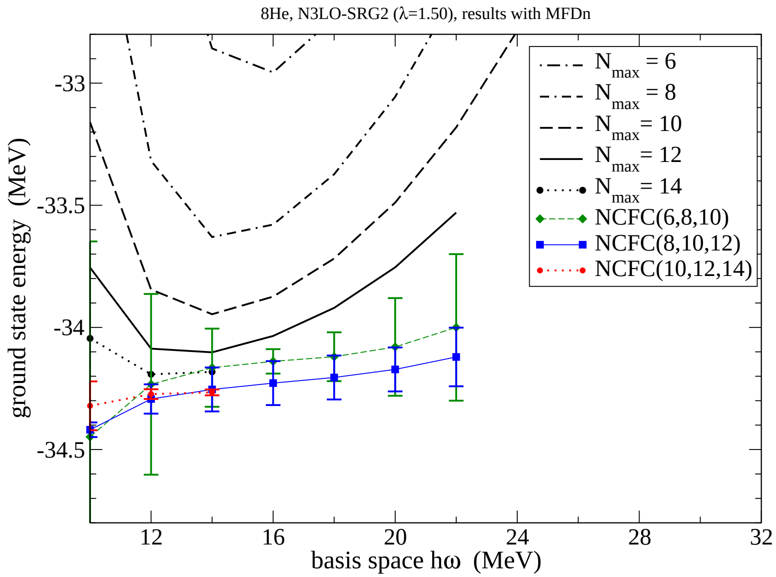

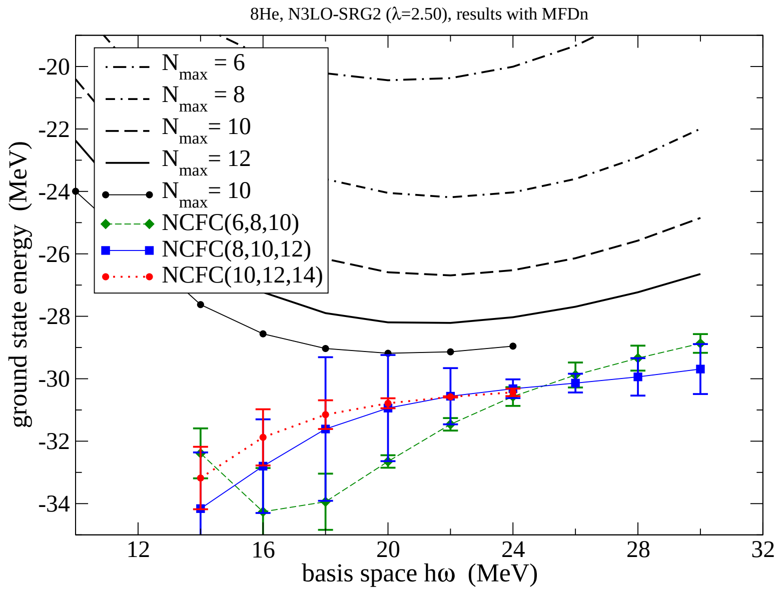

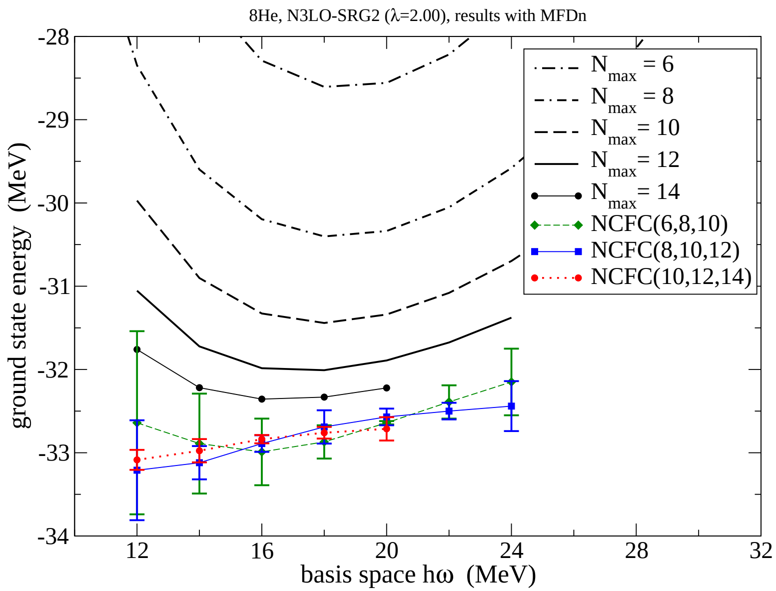

- Convergence pattern for $\lambda = 1.5$ and $\lambda = 2.5$ are qualitatively different
- Optimal $\hbar\omega$ shifts from $\hbar\omega = 14$ MeV to $\hbar\omega = 18$ MeV to $\hbar\omega = 22$ MeV

RMS radii of ^{8}He

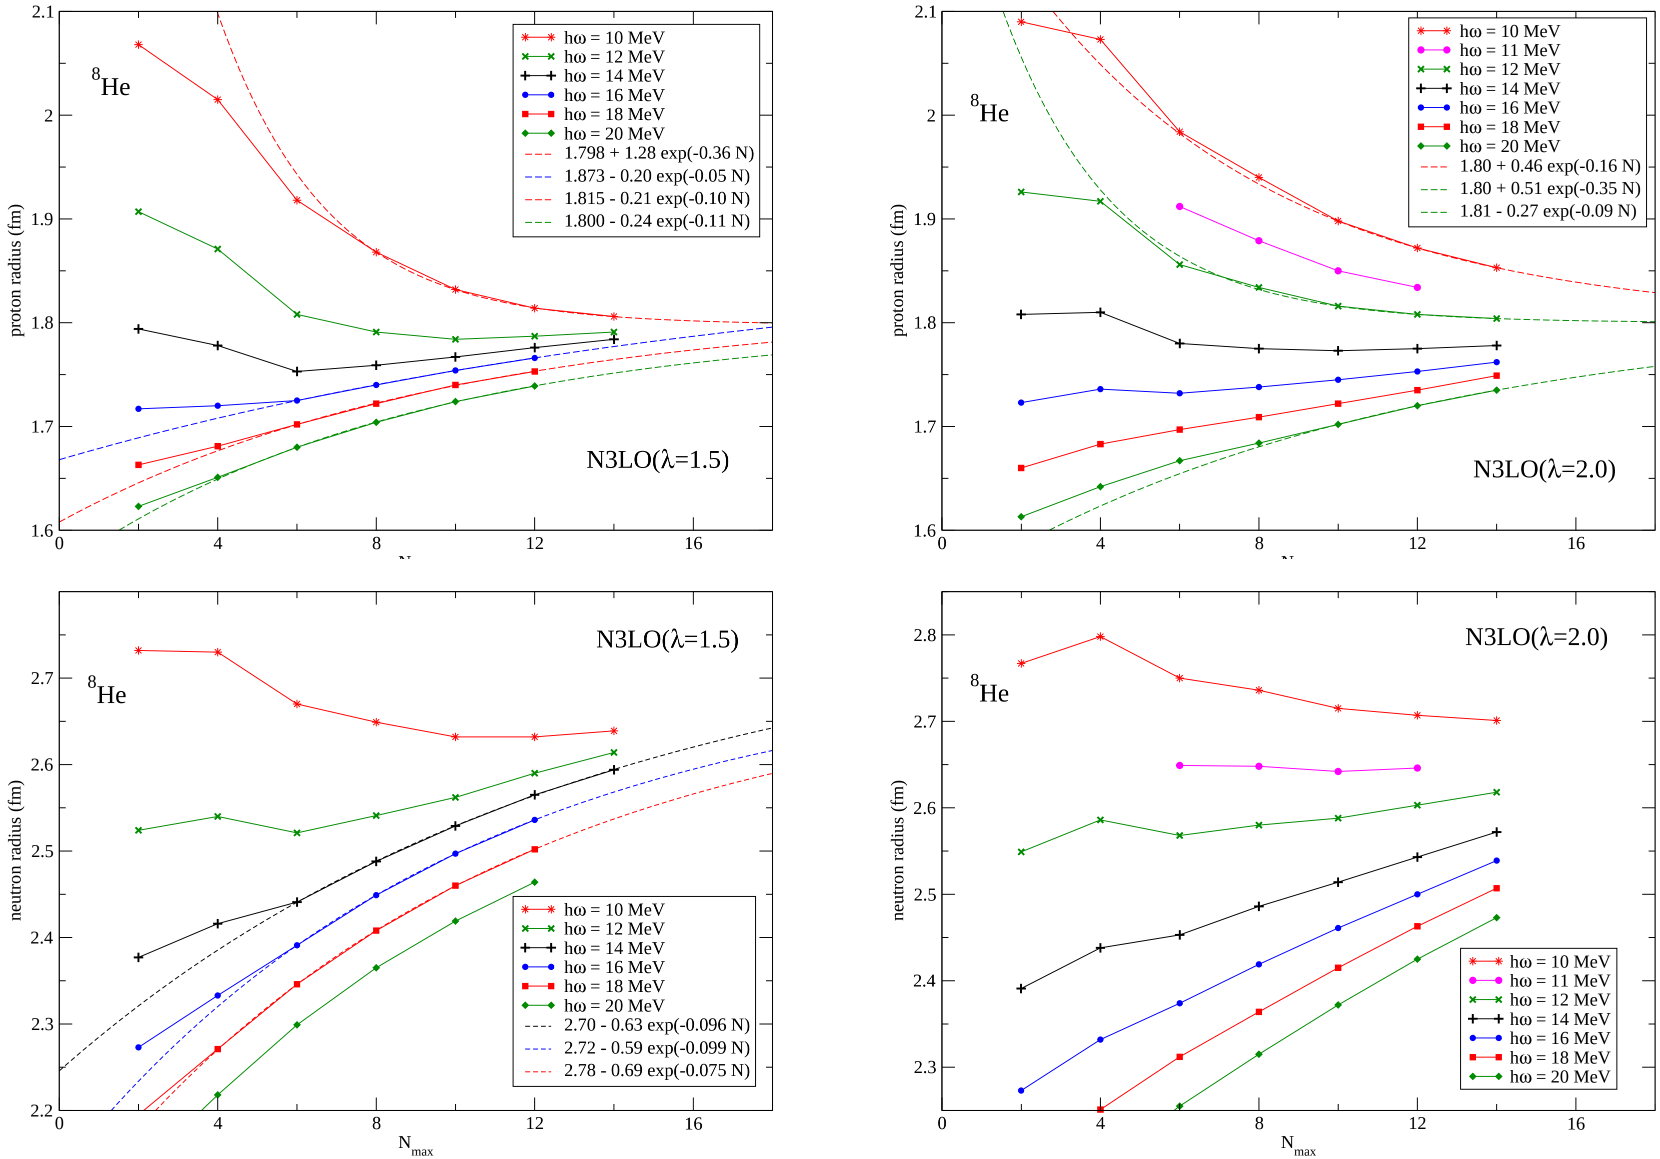

Proton and neutron densities ^{8}He

- Translationally-invariant, after deconvolution of Center-of-Mass

$$\rho_{\rm ti}(\vec{r}) \quad = \quad F\left[\frac{F[\rho^\omega(\vec{r})]}{F[\rho^\omega_{\rm cm}(\vec{r})]}\right]$$

Cockrell, Vary, PM, arXiv:1201.0724 [nucl-th]

- become independent of basis $\hbar\omega$ for large $N_{\max}$

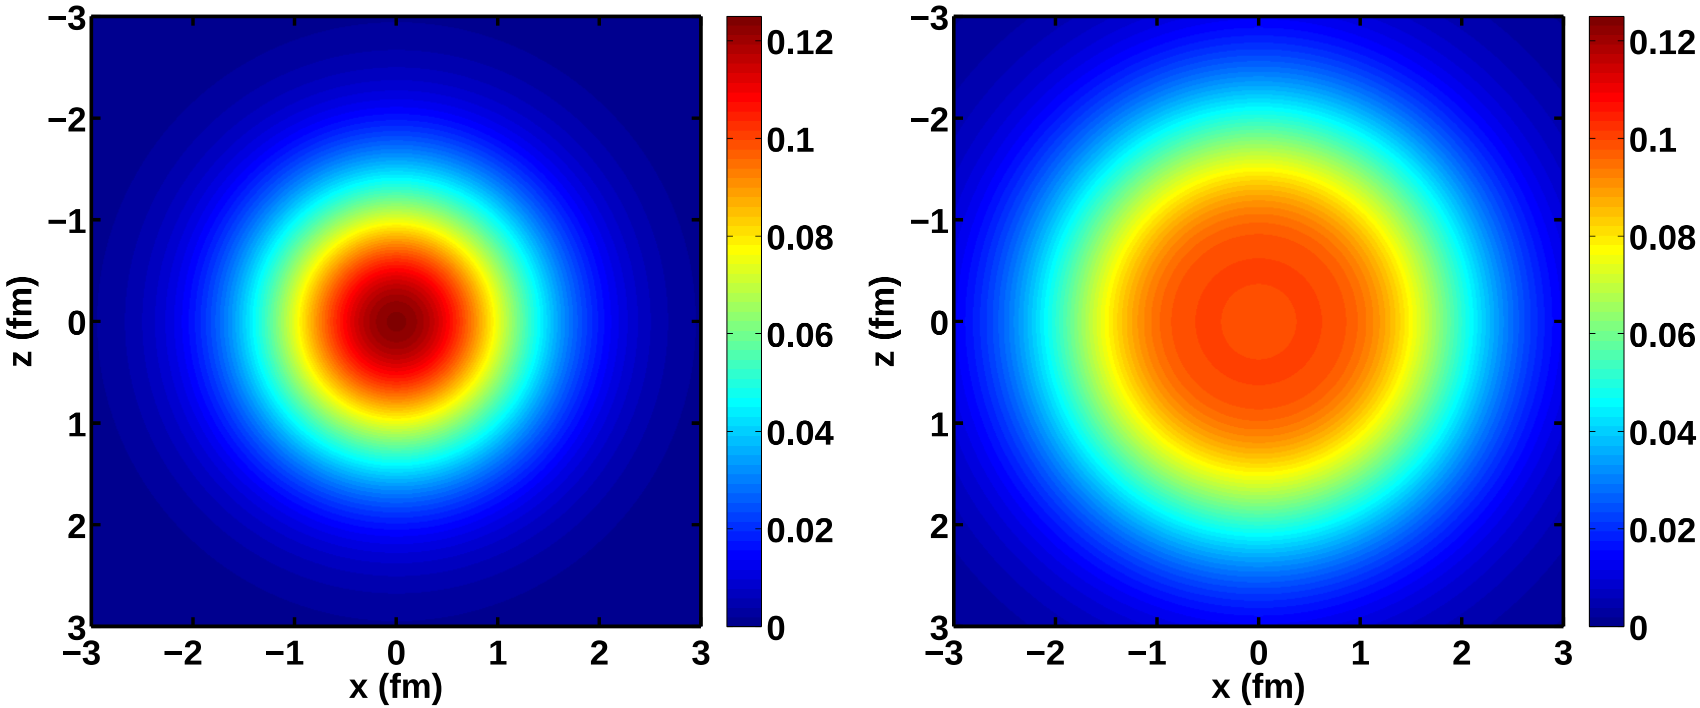

- plotted for $\lambda = 1.5$, $N_{\max} = 12$, $\hbar\omega = 12$ MeV

Ground state energy of ^{12}C

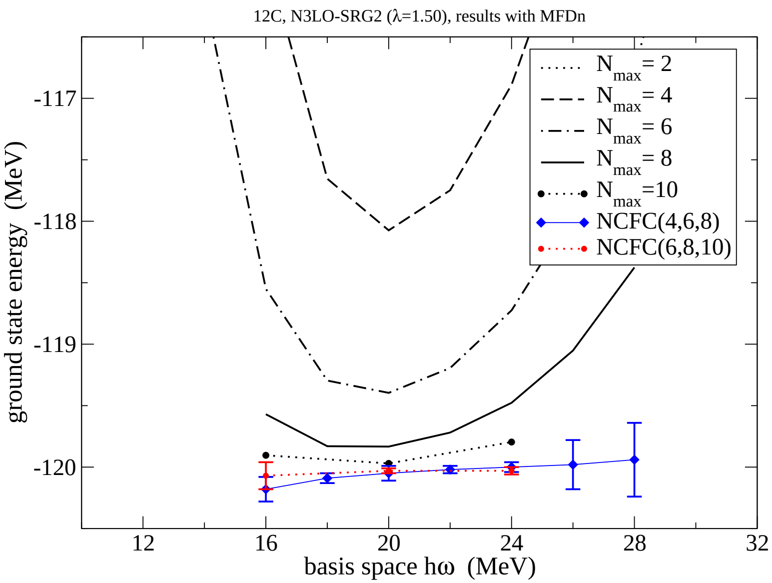

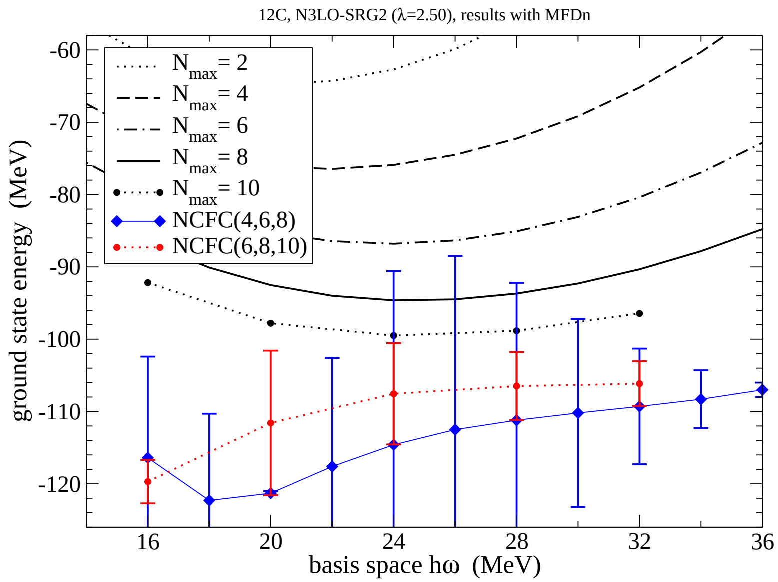

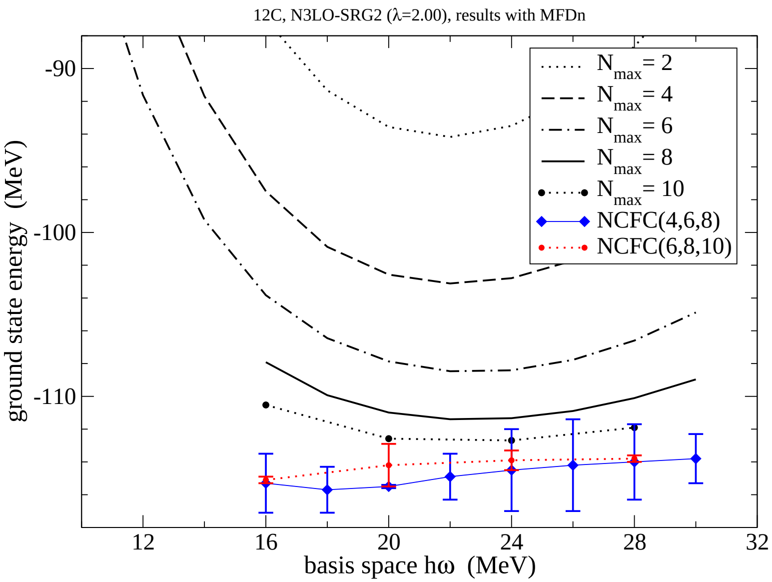

- Convergence pattern for $\lambda = 1.5$ and $\lambda = 2.5$ are qualitatively different
- Optimal $\hbar\omega$ shifts from $\hbar\omega = 22$ MeV to $\hbar\omega = 28$ MeV to $\hbar\omega = 32 \sim 36$ MeV

RMS radius of ^{12}C

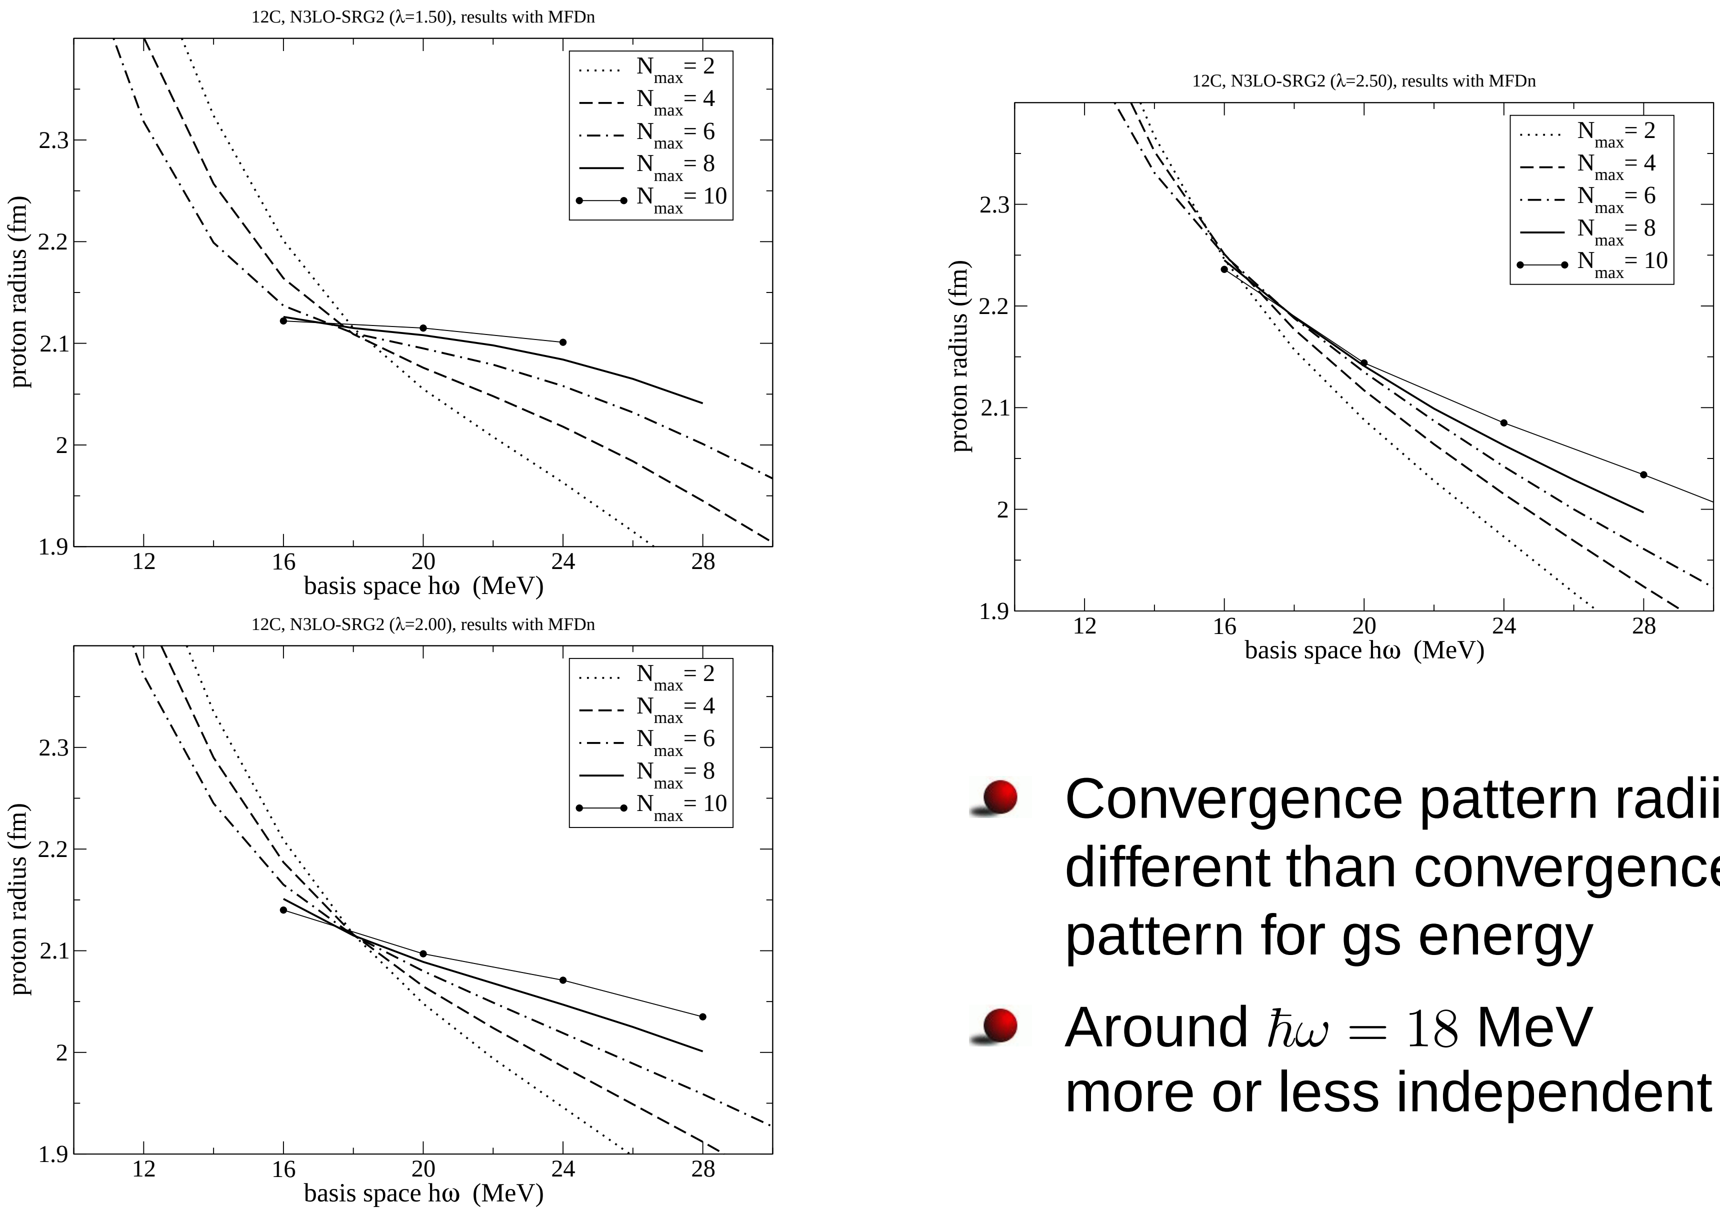

- Convergence pattern radii very different than convergence pattern for gs energy
- Around $\hbar\omega = 18$ MeV more or less independent of λ

Excitation energy 2^+ and 4^+ of ^{12}C

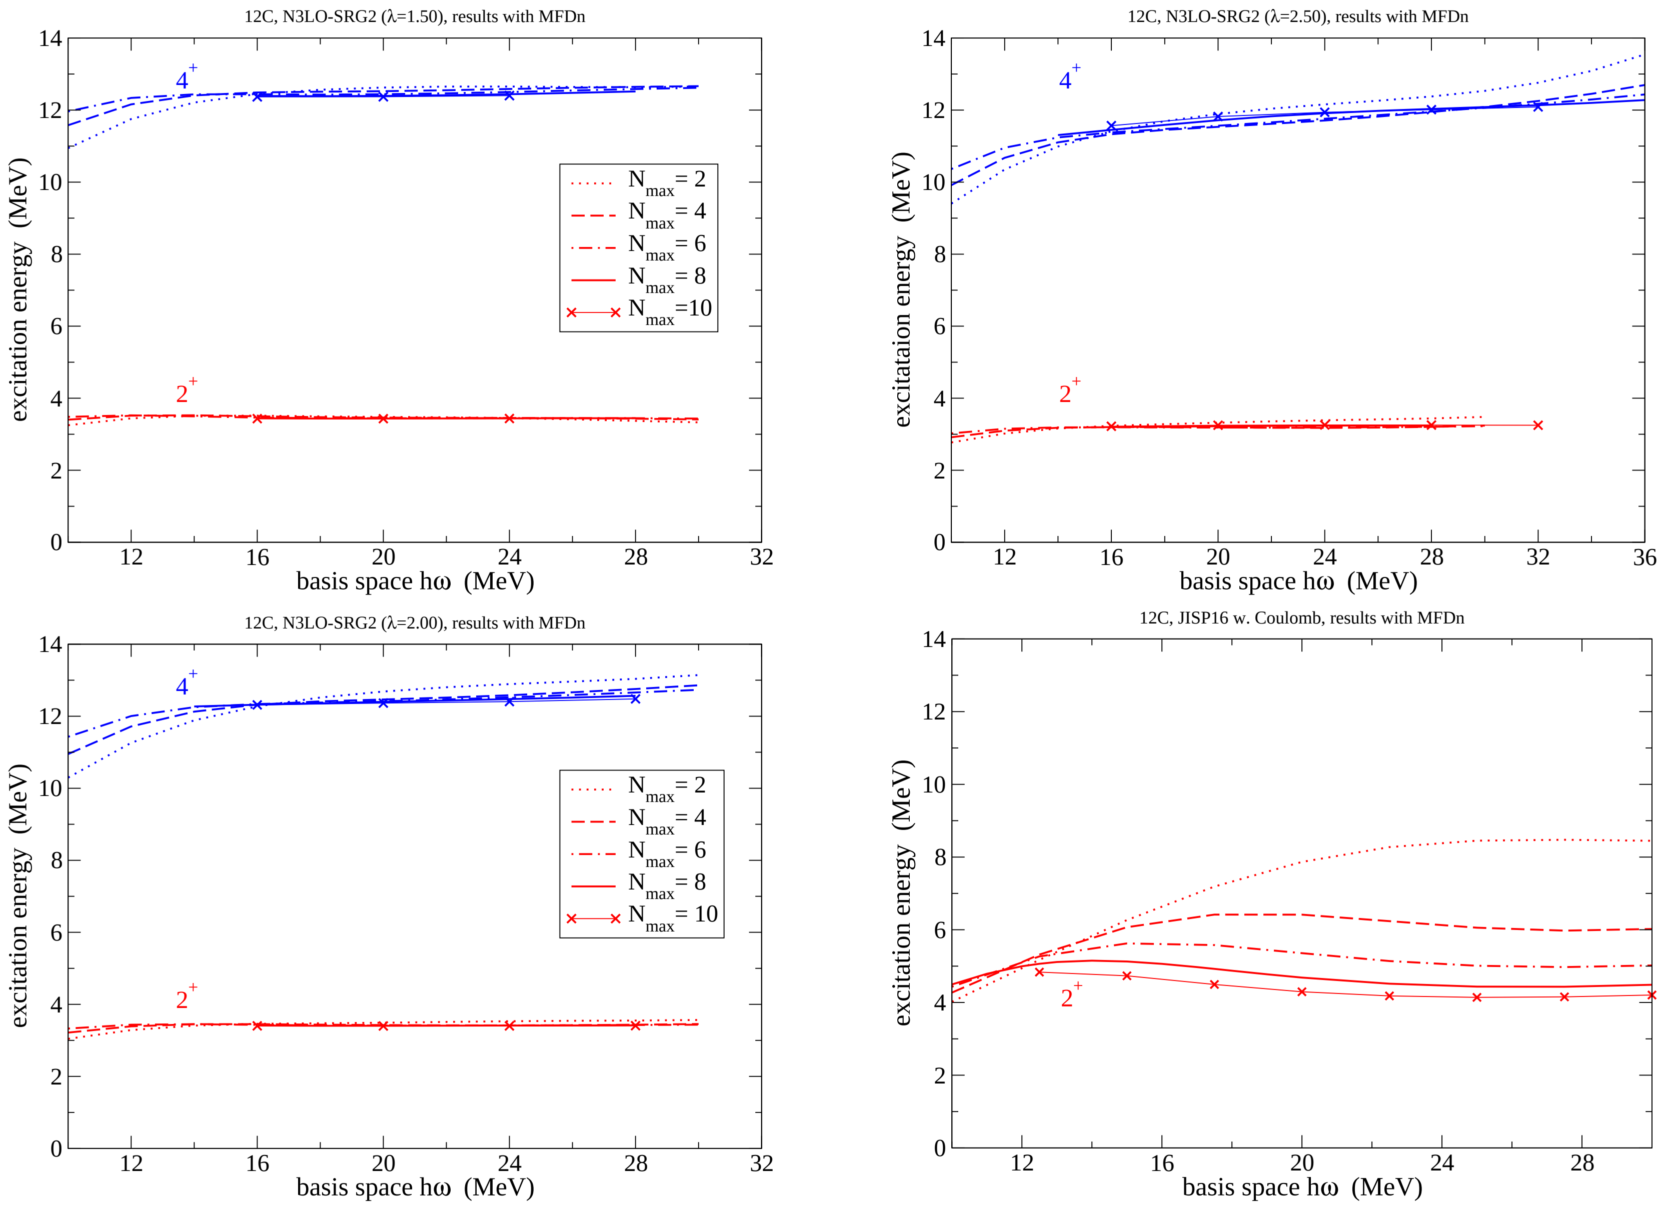

Quadrupole moment 2^+ of ^{12}C

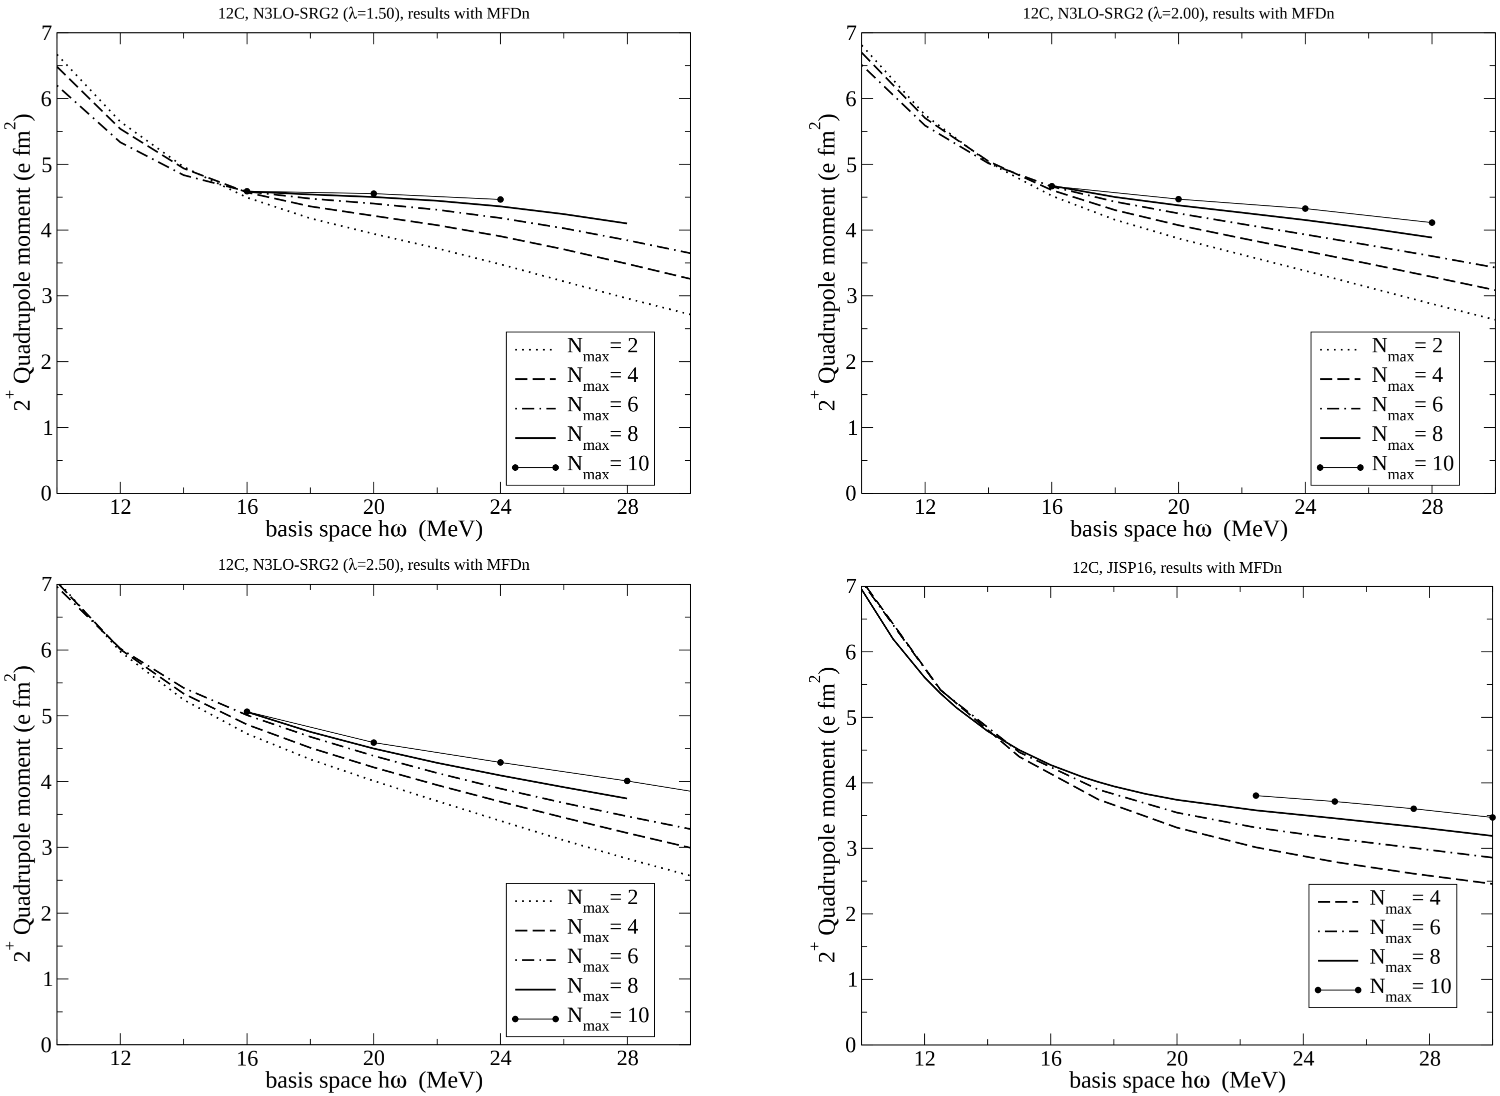

Translationally-invariant ^{12}C proton densities

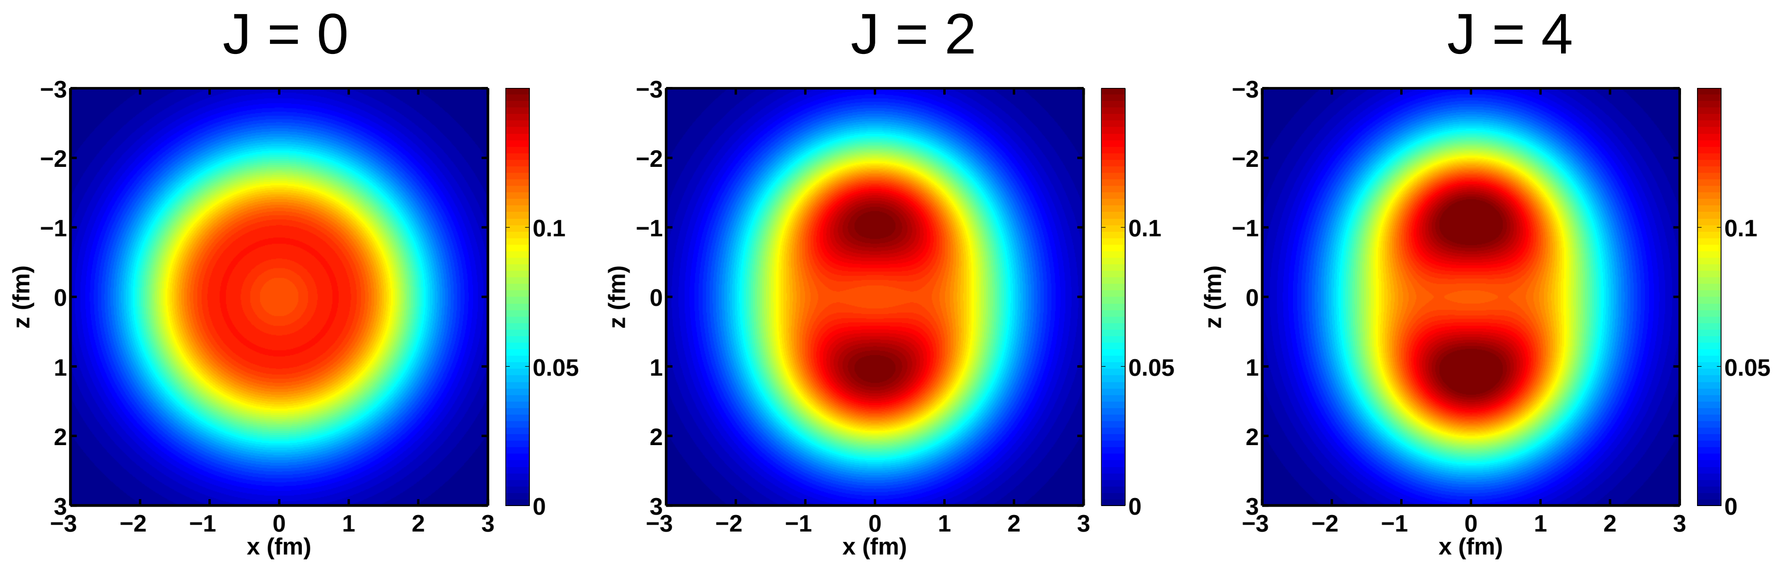

Translationally-invariant densities

- plotted for $M_J = J$
- plotted for $\lambda = 1.5$, $N_{\text{max}} = 10$, $\hbar\omega = 16$ MeV
- become independent of basis $\hbar\omega$ for large N_{max} after deconvolution of Center-of-Mass motion

Cockrell, PhD student

$$\rho_{\rm ti}(\vec{r}) \quad = \quad F\left[\frac{F[\rho^{\omega}(\vec{r})]}{F[\rho^{\omega}_{\rm cm}(\vec{r})]}\right]$$

Cockrell, Vary, PM, arXiv:1201.0724 [nucl-th]

Ground state energy of ^{16}O

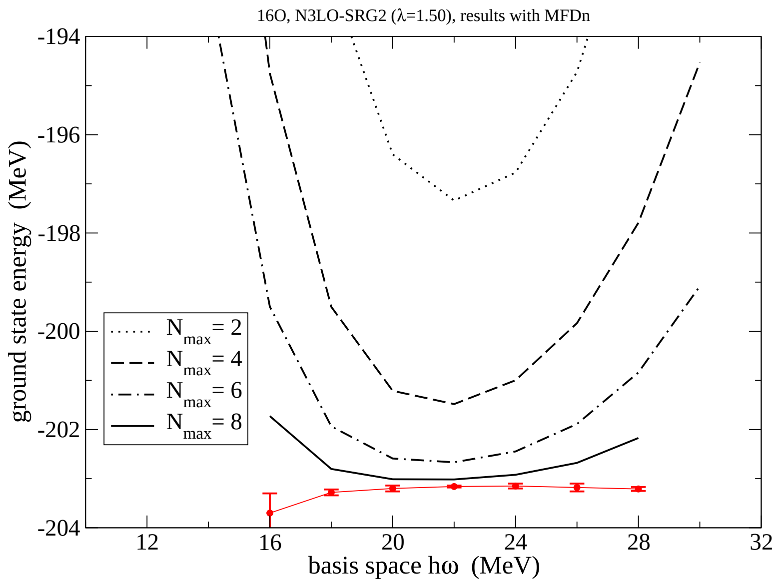

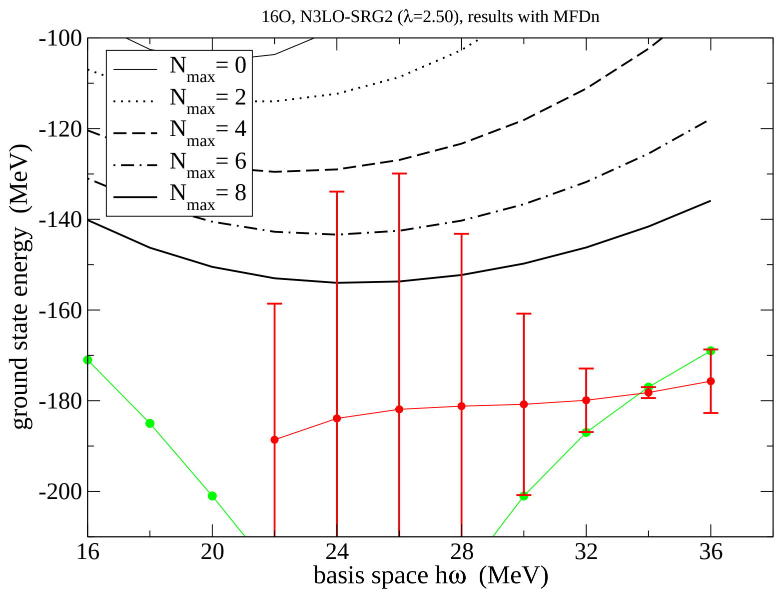

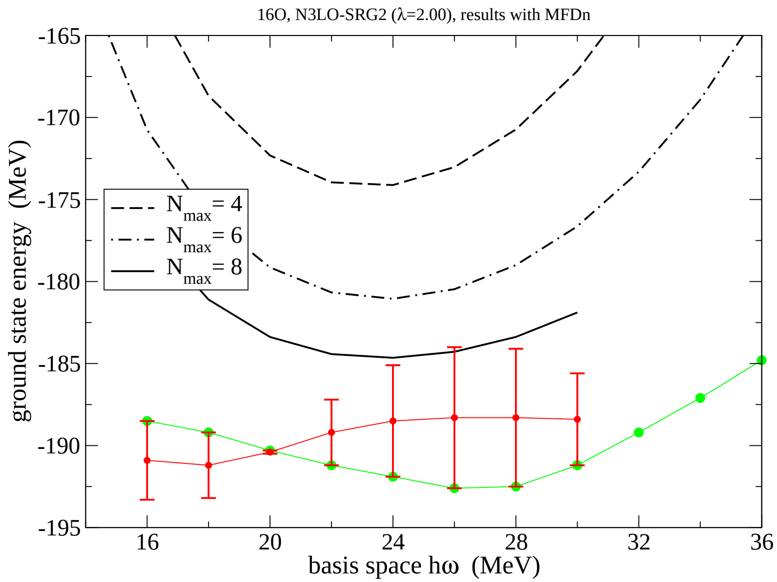

- Convergence pattern for $\lambda = 1.5$ and $\lambda = 2.5$ are qualitatively different
- Optimal $\hbar\omega$ shifts from $\hbar\omega = 22$ MeV to $\hbar\omega = 26 \sim 30$ MeV to $\hbar\omega = 30 \sim 36$ MeV

Ground state energy of ^{18}O

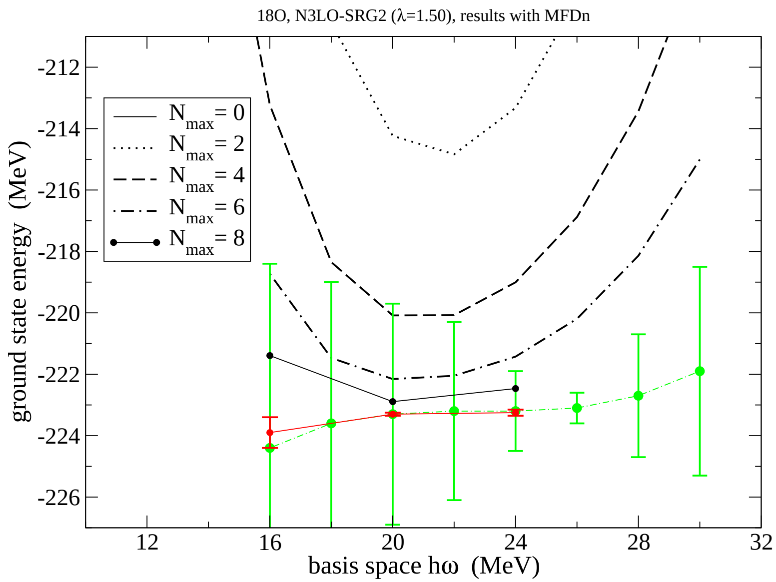

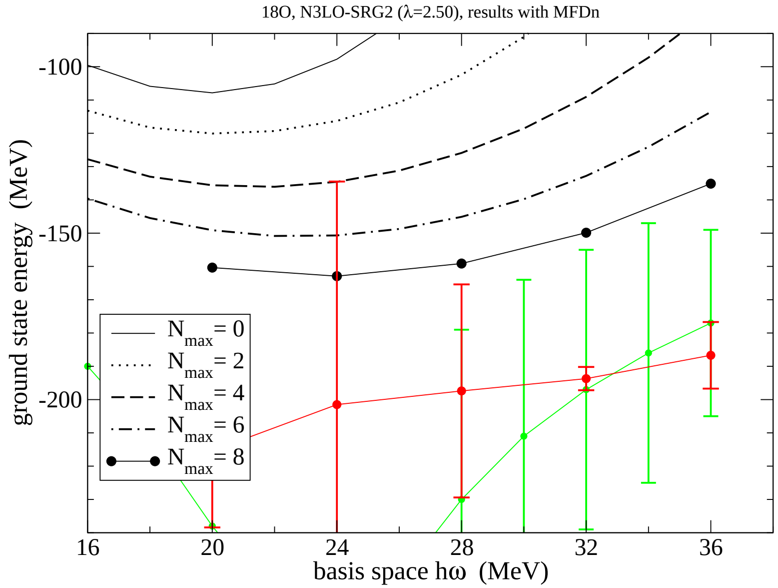

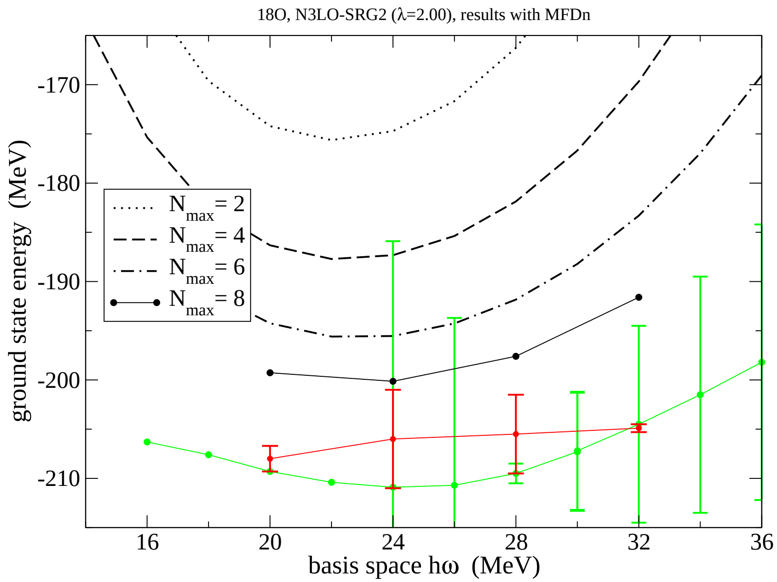

- Convergence pattern for $\lambda = 1.5$ and $\lambda = 2.5$ are qualitatively different
- Optimal $\hbar\omega$ shifts from $\hbar\omega = 22$ MeV to $\hbar\omega = 28 \sim 32$ MeV to $\hbar\omega = 30 \sim 36$ MeV

Translationally-invariant ^{18}O proton and neutron densities

gs 0^+ neutron neutron-proton proton

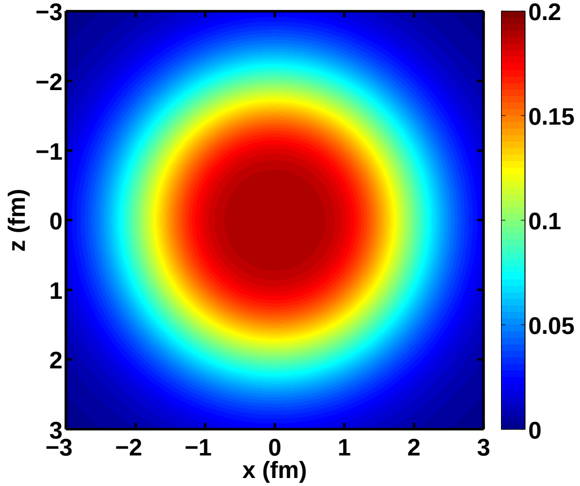

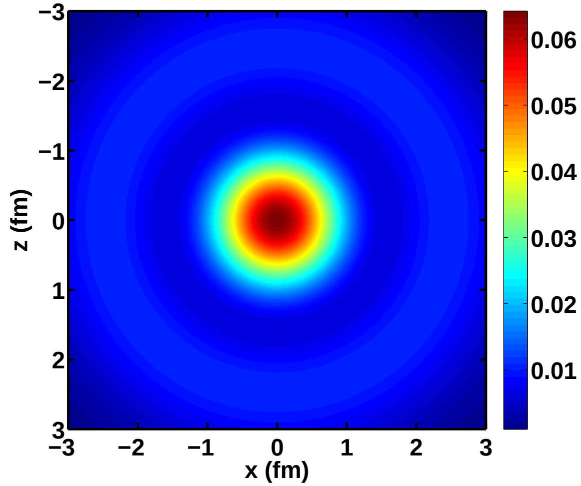

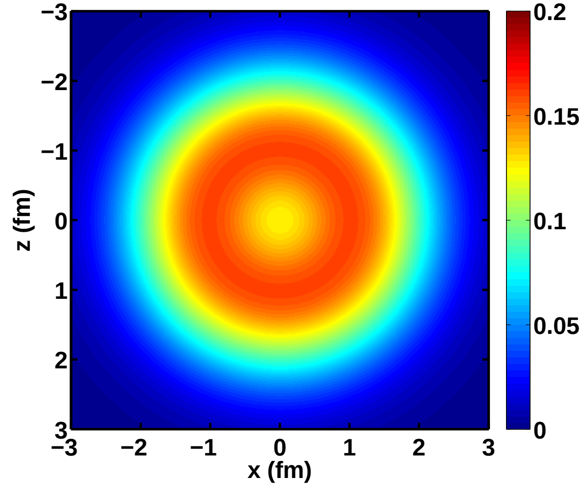

excited 2^+ neutron neutron-proton proton

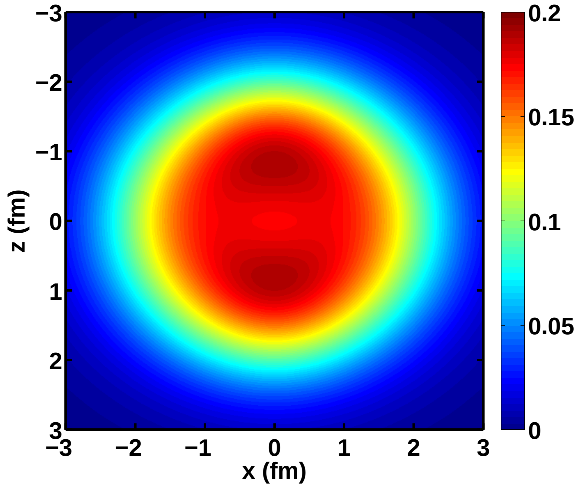

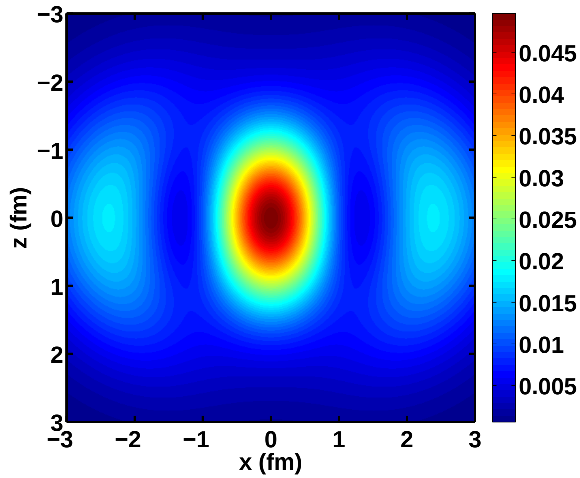

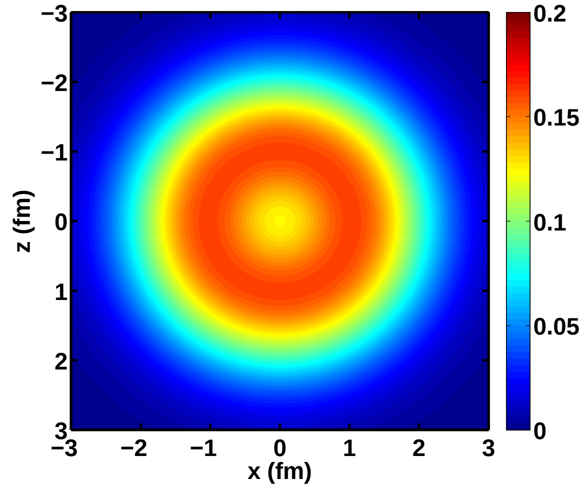

Inclusion of induced and explicit 3-body forces

- Consider SRG-evolved chiral interactions
  - chiral N3LO NN-only
  - chiral NN plus induced 3-body forces
  - chiral N3LO NN plus N2LO 3NF
- Use NN-only calculations up to $N_{\text{max}} = 10$ or 12 to validate extrapolation and establish optimal $\hbar\omega$ region
- Extrapolate 3-body calculations up to $N_{\text{max}} = 8$ assuming that optimal $\hbar\omega$ region for NN-only calculations is also (more or less) optimal for 3-body calculations
  - should be valid if there is a hierarchy for many-body forces

$$V_{NN} \gg V_{NNN} \gg V_{NNNN}$$

- Work in progress: ^{7}Li, ^{7}Be, ^{8}Be, ^{10}B, ^{12}C

w. Furnstahl, Jurgenson, Navratil, Ormand, Vary

Results for 3^+ state of ^{10}B at $\lambda = 1.0$

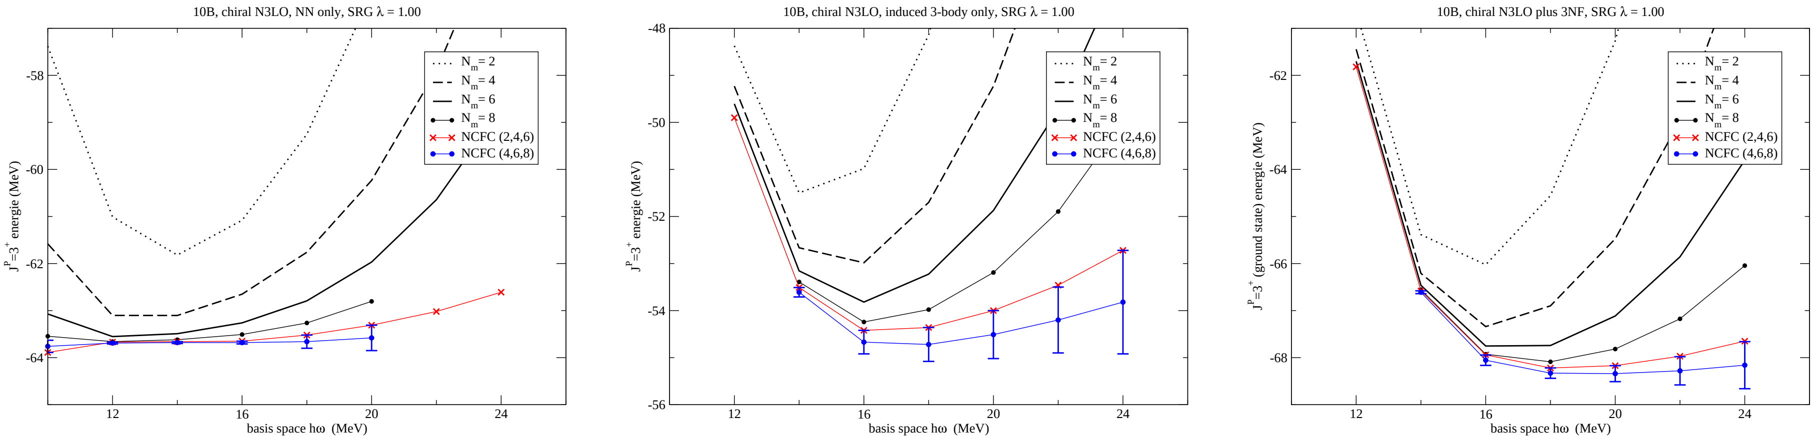

- NN-only calculations: optimal $\hbar\omega$ region is 12 to 16 MeV
- 3-body calculations in optimal $\hbar\omega$ region
  - rapid convergence with N_{max}
  - however 'converged' results are $\hbar\omega$ dependent because the truncation on the initial 3-body space, $N_{\text{max}} = 40$ is insufficient in this $\hbar\omega$ range?
- 3-body calculations restricted to $\hbar\omega \geq 20$ MeV
  - slower convergence with N_{max},
  - approximately independent of $\hbar\omega$

Results for 3^+ *state of* ^{10}*B at* $\lambda = 1.5$

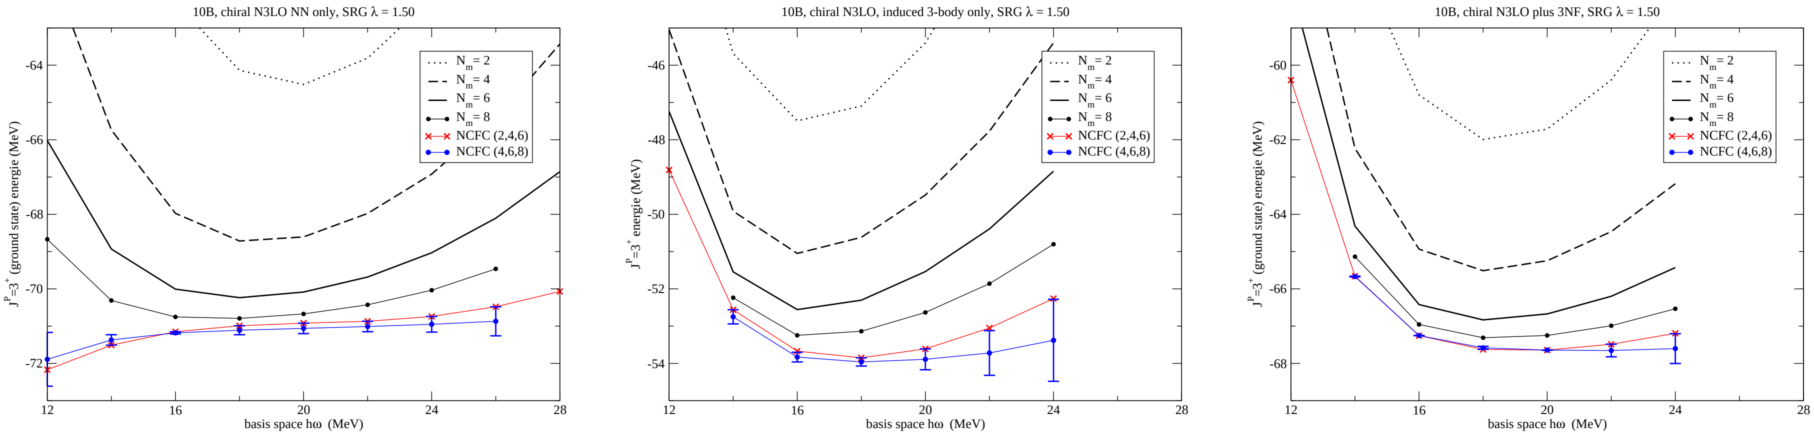

- NN-only calculations: optimal $\hbar\omega$ region is 16 to 22 MeV
- 3-body calculations in optimal $\hbar\omega$ region
  - rapid convergence with N_{max}
  - however 'converged' results with 3NF are $\hbar\omega$ dependent because the truncation on the initial 3-body space, $N_{\text{max}} = 40$ is insufficient in this $\hbar\omega$ range?
- 3-body calculations restricted to $\hbar\omega \geq 20$ MeV
  - slower convergence with N_{max},
  - approximately independent of $\hbar\omega$

Results for 3^+ state of ^{10}B at $\lambda = 2.0$ and $\lambda = 2.5$

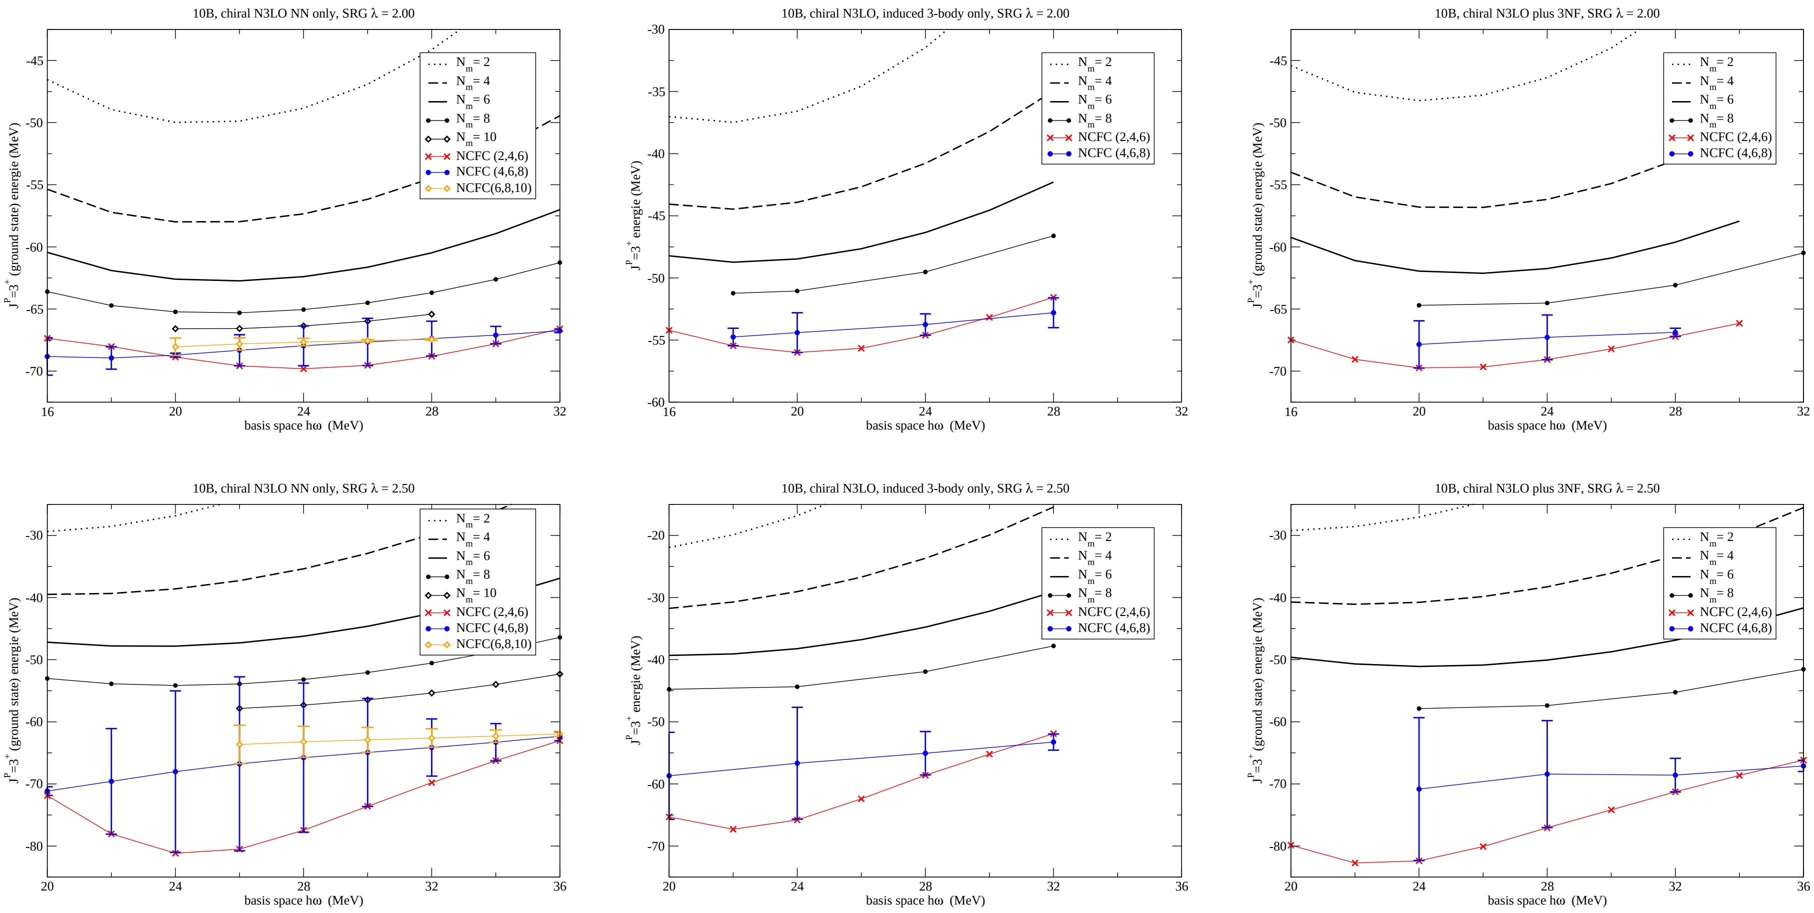

- NN-only calculations up through $N_{\text{max}} = 10$
  - optimal $\hbar\omega$ region is 24 to 30 MeV for $\lambda = 2.0$
  - optimal $\hbar\omega$ region is 30 to 40 MeV for $\lambda = 2.5$

Evidence for induced 4-body forces?

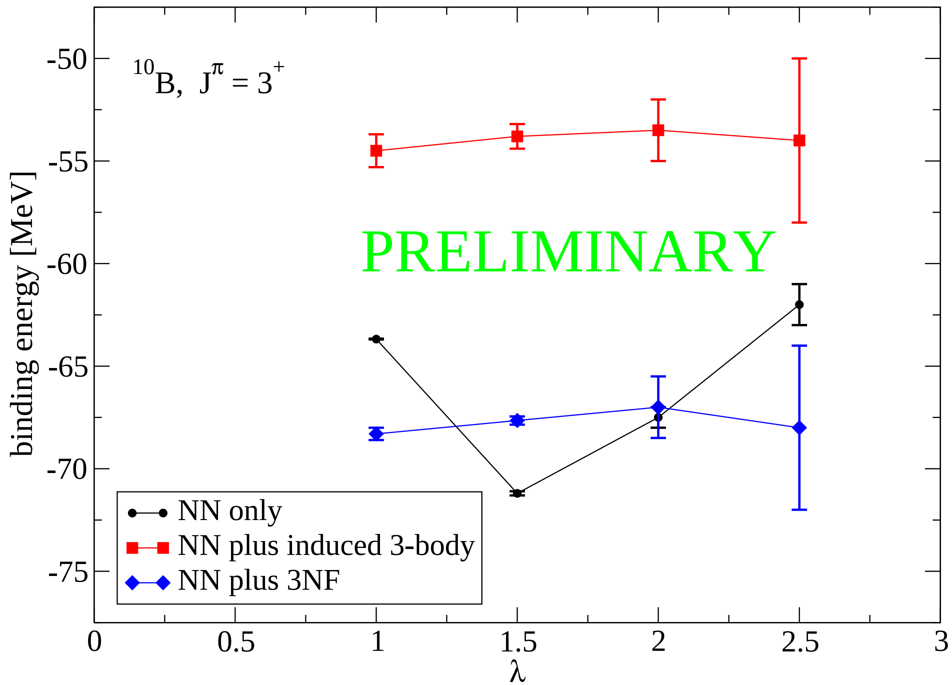

- Error estimates at $\lambda = 2.0$ and $\lambda = 2.5$ dominated by extrapolation error of many-body calculation
- Error estimates at $\lambda = 1.0$ and $\lambda = 1.5$ combination of numerical uncertainty in input 3-body matrix elements and of extrapolation

Conclusions

- MFDn: Scalable and load-balanced CI code for nuclear structure
  - new version under development, has run on 200k+ cores on Jaguar (ORNL) enabling largest model-space calculations
- Significant benefits from collaboration between nuclear physicists, applied mathematicians, and computer scientists
- Need at least $N_{\text{max}} = 8$ for reliable extrapolation to infinite basis
  - main challenge: construction and diagonalization of extremely large (D > 1 billion) sparse matrices
- Accurate results, including numerical error estimates, for p-shell nuclei with SRG evolved chiral NN-only interactions
- Calculations and extrapolations for p-shell nuclei with SRG evolved chiral 3-body forces in progress
  - challenge:
    complete $N_{\text{max}} = 10$ calculations with 3-body forces

Accuracy of SRG evolved 3-body matrix elements?

- A3max ramp used here: 40, 38, 36, 34, 32, 30, 28, 26, 24, 20, . . .
- A3max ramp Darmstadt: 40, 40, 40, 36, 32, 28, 24, 24, 24, 24, . . .

Differences in binding energies between our ramp and Darmstadt ramp

	^{7}Li, $\frac{3}{2}^-$			^{10}B, 3^+			^{12}C, 0^+		
$N_{\rm max}$	16	20	24	16	20	24	16	20	24
2	.038	.008	.004	.160	.038	.017	.267	.067	.031
4	.042	.008	.004	.176	.039	.017	.294	.069	.030
6	.044	.008	.004	.190	.046	.025	.321	.079	.041
8	.048	.013	.009		.061	.041			

- Note strong increase with number of nucleons
- Note systematic increase with $N_{\rm max}$

Need systematic investigation of convergence
of 3-body matrix elements as function of A3max