Microscopic Nuclear Structure Theory – Case Study

James P. Vary, Iowa State University

NERSC Workshop
April 29-30, 2014

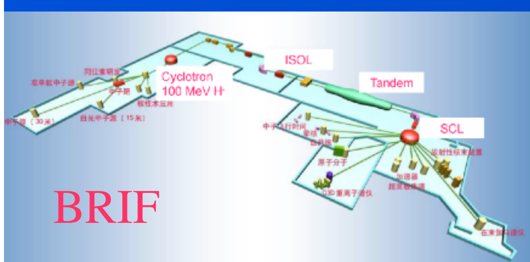

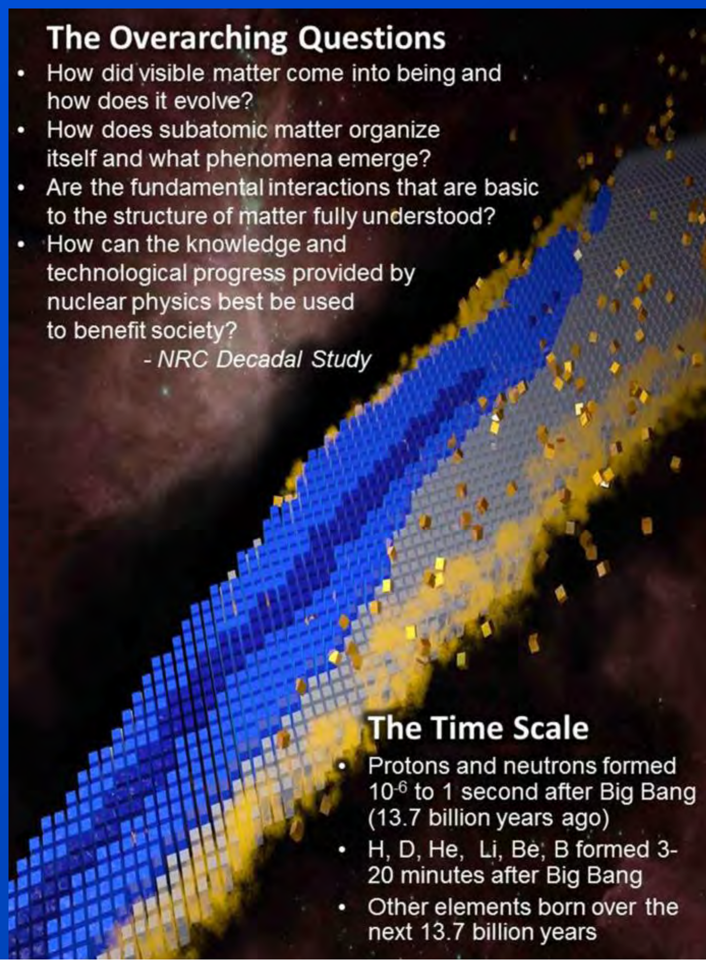

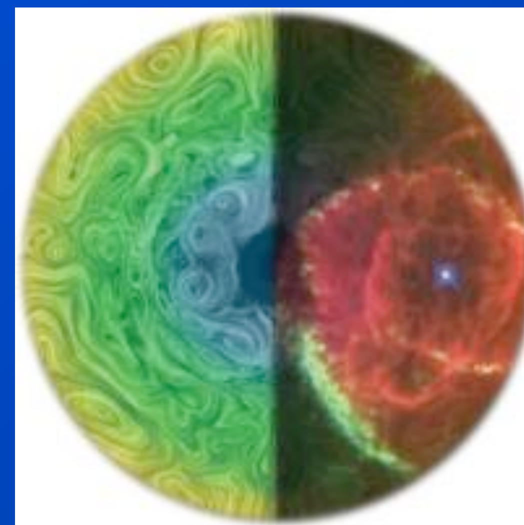

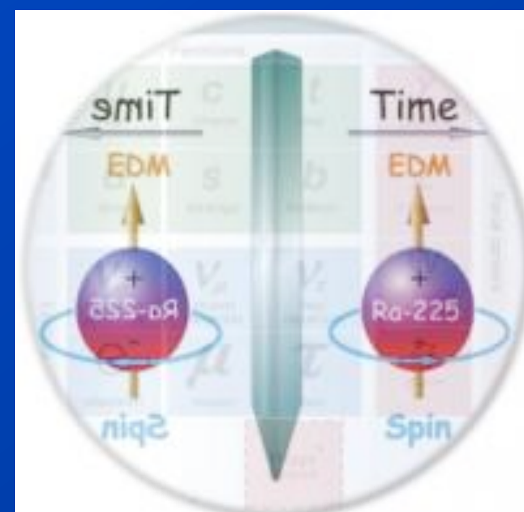

Overarching Problem

Main hypothesis

If the Standard Model is correct, we should be able
to accurately describe all nuclear processes

Long-term goal

Use all fundamental interactions
including yet-to-be-discovered interactions
to construct a model
for the evolution of the entire universe

Requirements

Major progress with basic theory, algorithms
and supercomputer simulations

Fundamental questions of nuclear physics => discovery potential

- What controls nuclear saturation?
- How shell and collective properties emerge from the underlying theory?
- What are the properties of nuclei with extreme neutron/proton ratios?
- Can we predict useful cross sections that cannot be measured?
- Can nuclei provide precision tests of the fundamental laws of nature?
- Can we solve QCD to describe hadronic structures and interactions?

Edison

U.S. DEPARTMENT OF ENERGY | Office of Science

National Science Foundation
WHERE DISCOVERIES BEGIN

SciDAC
Scientific Discovery through Advanced Computing

INCITE
Advancing America's Science and Industrial Competitiveness

+ K-super.
+ Blue Waters
+ TianHe II
+ Tachyon-II

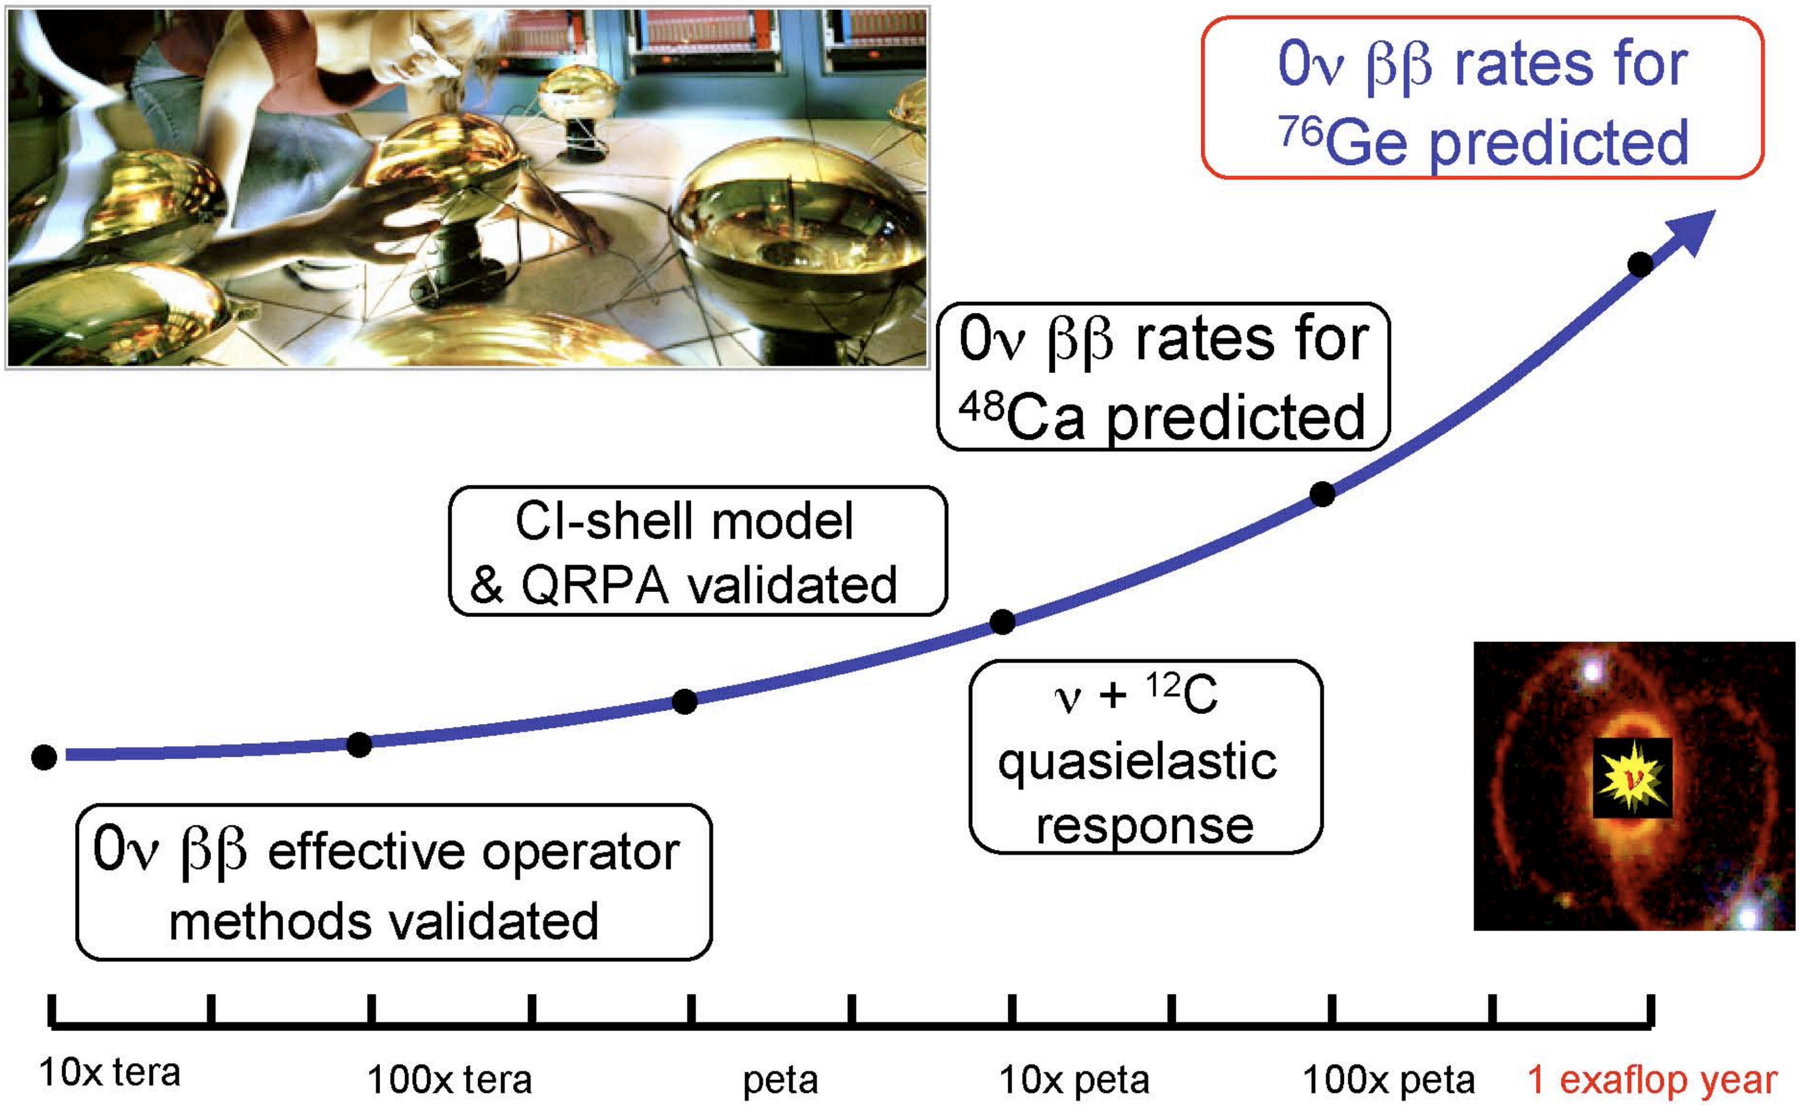

http://extremecomputing.labworks.org/nuclearphysics/report.stm

SciDAC Scientific Discovery through Advanced Computing

NUclear Computational Low-Energy Initiative

(2012-2017)

fusion

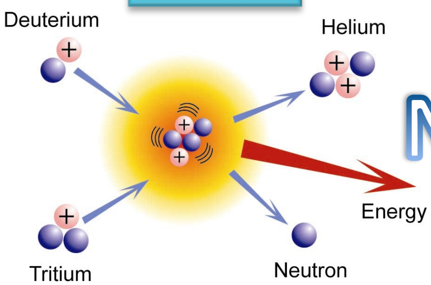

NIF

TJNAF

LENP facilities

FRIB

Validated Nuclear Interactions

Chiral EFT
Ab-initio

Optimization
Model validation
Uncertainty Quantification

Structure and Reactions: Light and Medium Nuclei

Neutron drops
EOS
Correlations

Structure and Reactions: Heavy Nuclei

Stellar burning

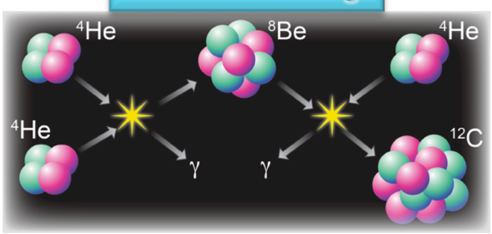

Ab-initio
RGM
CI

Load balancing
Eigensolvers
Nonlinear solvers
Model validation
Uncertainty Quantification

DFT
TDDFT

Load balancing
Optimization
Model validation
Uncertainty Quantification
Eigensolvers
Nonlinear solvers
Multiresolution analysis

Neutrinos and Fundamental Symmetries

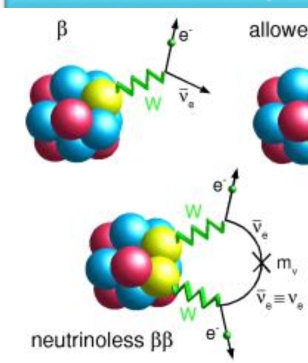

SNS
Majorana

Neutron Stars

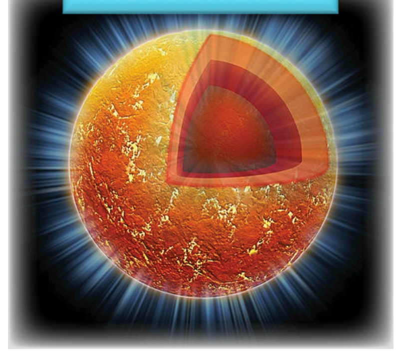

Fission

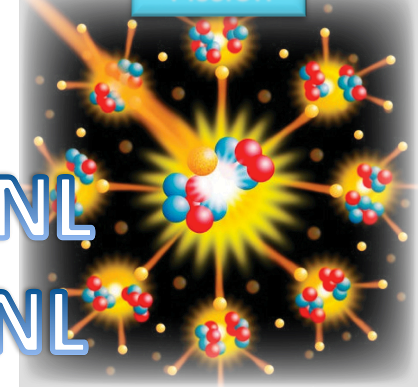

LANL
LLNL

computingnuclei.org (adapted by Gaute Hagen)

Next Generation Nuclear Hamiltonians:

Pionful Chiral EFT = Fully consistent through N3LO (LENPIC)
Deltaful Chiral EFT = R. Machleidt
Strangeful Chiral EFT = ?
Veff from LQCD = NPLQCD (M. Savage, et al)

Each puts major demands on the many-body theory

Growing demands => larger collaborating teams, growing computational resources,
Increase in the multi-disciplinary character (SciDAC), . . .

Calculation of three-body forces at N^3LO

Low
Energy
Nuclear
Physics
International
Collaboration

J. Golak, R. Skibinski, K. Tolponicki, H. Witala

RUB — E. Epelbaum, H. Krebs

JÜLICH — A. Nogga

The Ohio State University — R. Furnstahl

Technische Universität Darmstadt — S. Binder, A. Calci, K. Hebeler, J. Langhammer, R. Roth

P. Maris, J. Vary

Kyutech — H. Kamada

Goal

Calculate matrix elements of 3NF in a partial-wave decomposed form which is suitable for different few- and many-body frameworks

Challenge

Due to the large number of matrix elements, the calculation is extremely expensive.

Strategy

Develop an efficient code which allows to treat arbitrary local 3N interactions.
(Krebs and Hebeler)

The Nuclear Many-Body Problem

The many-body Schroedinger equation for bound states consists of $2\binom{A}{Z}$ coupled second-order differential equations in 3A coordinates using strong (NN & NNN) and electromagnetic interactions.

Successful *ab initio* quantum many-body approaches (A > 6)

Stochastic approach in coordinate space
Greens Function Monte Carlo (**GFMC**)

Meson Exchg interactions

Featured results here

Hamiltonian matrix in basis function space
No Core Configuration Interaction (**NCSM/NCFC**)

Cluster hierarchy in basis function space
Coupled Cluster (**CC**)

Lattice Nuclear Chiral EFT, MB Greens Function,
MB Perturbation Theory, . . . approaches

Chiral EFT interactions

Comments

All work to preserve and exploit symmetries
Extensions of each to scattering/reactions are well-underway
They have different advantages and limitations

No Core Shell Model
A large sparse matrix eigenvalue problem

$$H = T_{rel} + V_{NN} + V_{3N} + \bullet\bullet\bullet$$

$$H\left|\Psi_i\right\rangle = E_i\left|\Psi_i\right\rangle$$

$$\left|\Psi_i\right\rangle = \sum_{n=0}^{\infty} A_n^i \left|\Phi_n\right\rangle$$

$$\text{Diagonalize } \left\{\left\langle\Phi_m\right|H\left|\Phi_n\right\rangle\right\}$$

- Adopt realistic NN (and NNN) interaction(s) & renormalize as needed - retain induced many-body interactions: Chiral EFT interactions and JISP16
- Adopt the 3-D Harmonic Oscillator (HO) for the single-nucleon basis states, α, β,…
- Evaluate the nuclear Hamiltonian, H, in basis space of HO (Slater) determinants (manages the bookkeeping of anti-symmetrization)
- Diagonalize this sparse many-body H in its "m-scheme" basis where [α =(n,l,j,m_j,τ_z)]

$$\left|\Phi_n\right\rangle = [a_\alpha^+ \bullet\bullet\bullet a_\varsigma^+]_n \left|0\right\rangle$$

$$n = 1,2,\ldots,10^{10} \text{ or more!}$$

- Evaluate observables and compare with experiment

Comments

- Straightforward but computationally demanding => new algorithms/computers
- Requires convergence assessments and extrapolation tools
- Achievable for nuclei up to A=16 (40) today with largest computers available

Ab initio NCSM Flowcharts

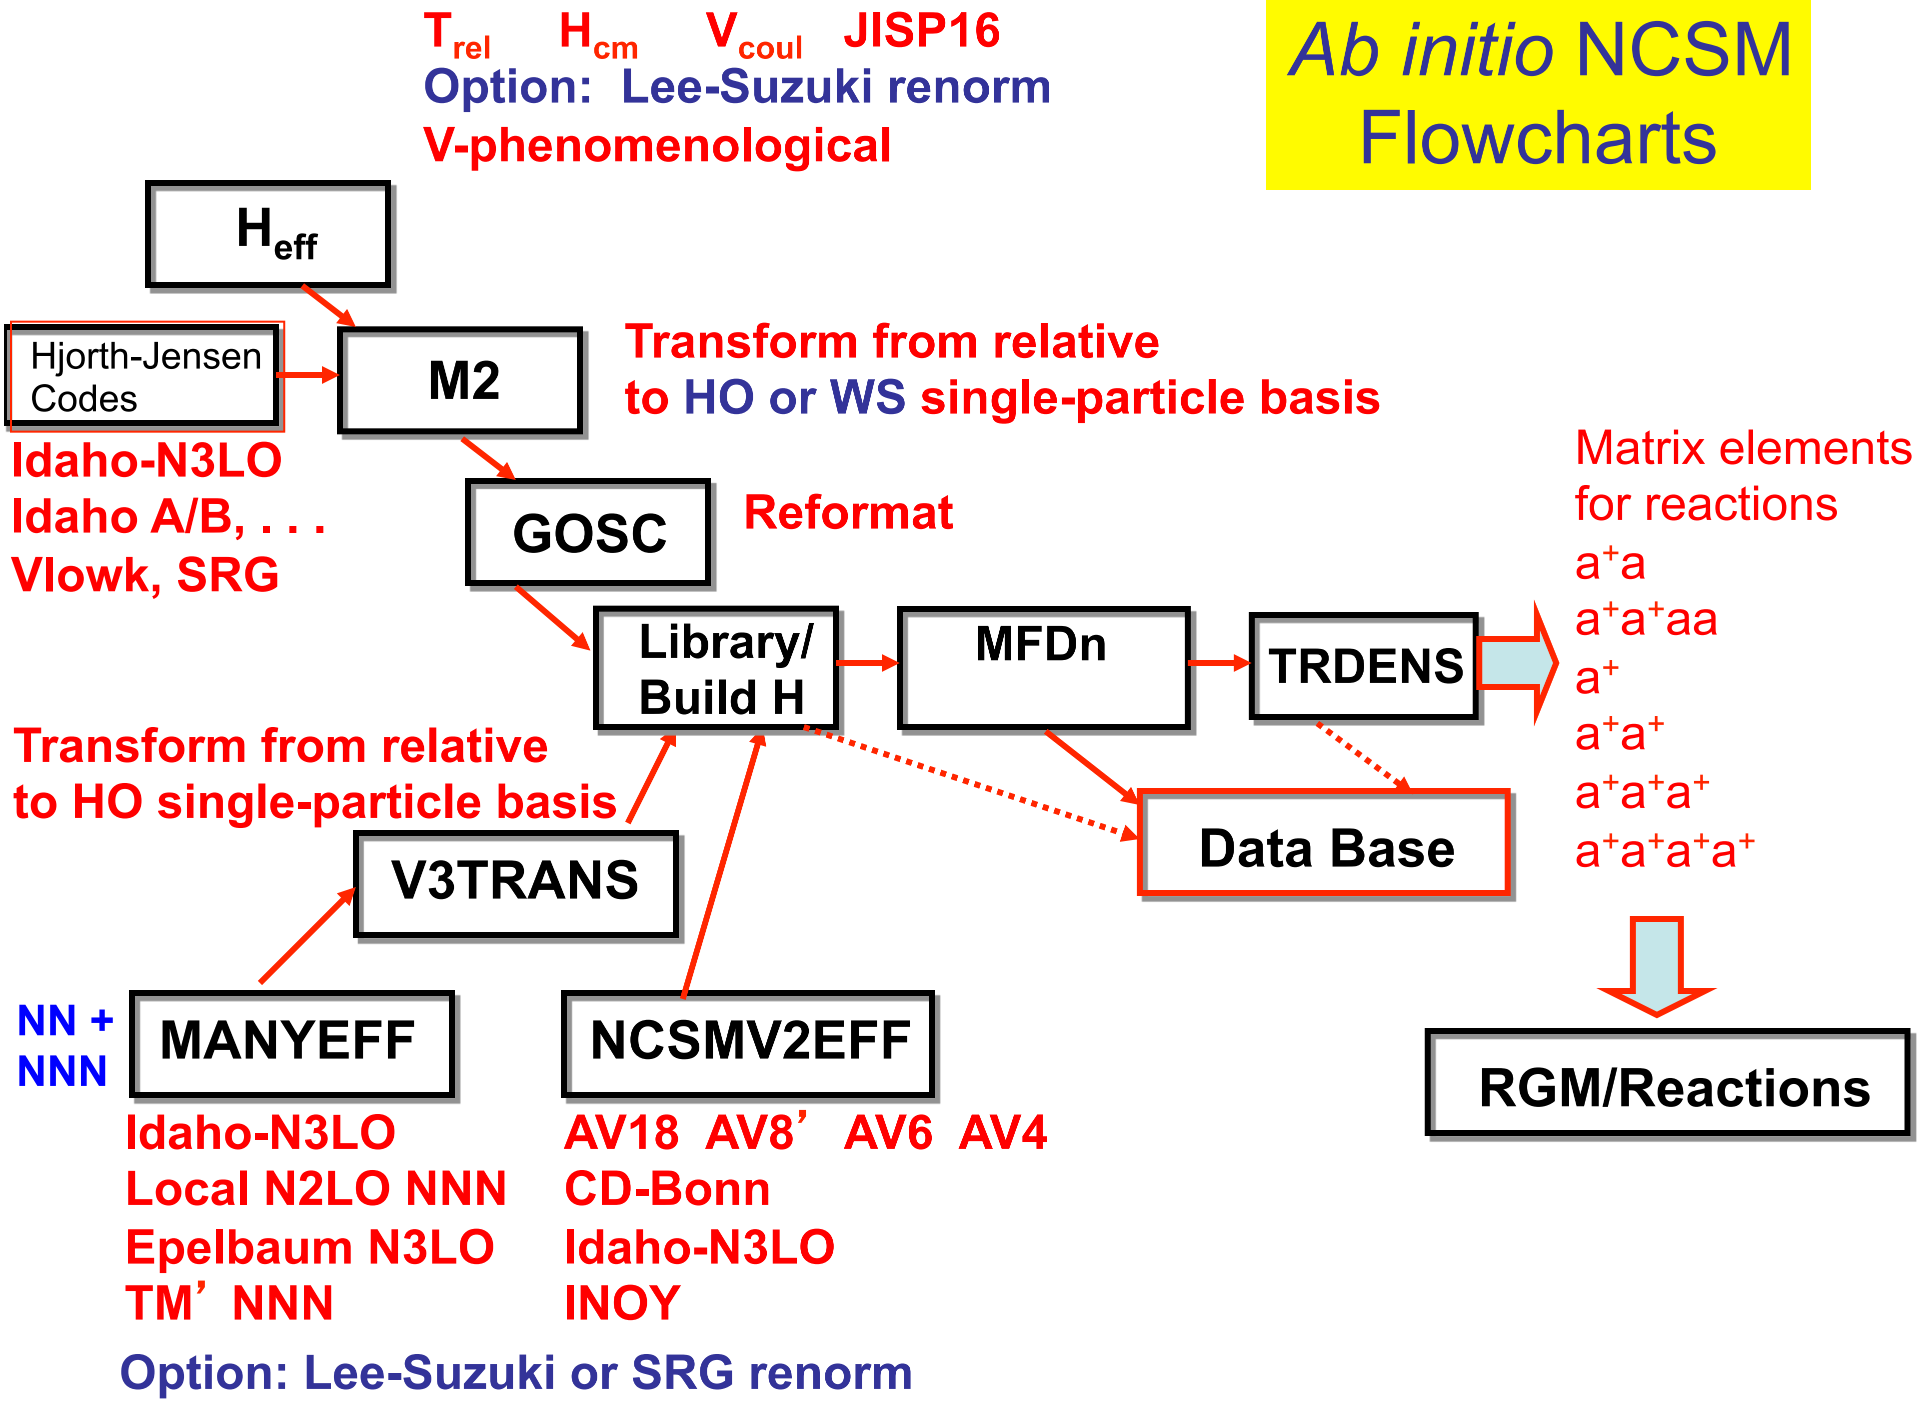

“Many-Fermion Dynamics – nuclear” or “MFDn”

The primary code we have developed and continue to improve

There are 7 major stages of the calculations performed by MFDn:

1. Enumerate the many-body basis space according to user-defined criteria
2. Determine the location of non-zero many-body matrix elements in this basis space and hence the number of non-zeroes that must be evaluated for the full Hamiltonian
3. Read in the nucleon-nucleon plus three-nucleon interaction files that define the Hamiltonian
4. Construct and store (partially or fully) the many-body Hamiltonian matrix
5. Perform the Lanczos iterations until either a fixed number of iterations are achieved or a convergence criterium is met. Perform orthonormalization of the Lanczos basis vectors after each iteration.
6. Transform the eigenvectors from the Lanczos basis back to the original basis
7. Use a selected set of the eigenvectors in the original basis to calculate a suite of experimental observables and 1-body density matrices that, optionally, may be stored for reuse later. The option to evaluate and store 2-body density matrices is in the planning stage.

Low Energy NP Application Areas

Application	Production Run Sizes	Resource	Dense Linear Alg.	Sparse Linear Alg.	Monte Carlo
AGFMC: Argonne Green's Function Monte Carlo	262,144 cores @ 10 hrs	Mira			X
MFDn: Many Fermion Dynamics - nuclear	260K cores @ 4 hrs 500K cores @ 1.33 hrs	Titan Mira		X	
NUCCOR: Nuclear Coupled-Cluster Oak Ridge, m-scheme & spherical	100K cores @ 5 hrs (1 nucleus, multiple parameters)	Titan		X	
DFT Code Suite: Density Functional Theory, mean-field methods	100K cores @ 10 hrs (entire mass table, fission barriers)	Titan	X		
MADNESS: Schroedinger, Lippman-Schwinger and DFT	40,000 cores @ 12 hrs (extreme asymmetric functions)	Titan	X	X	
NCSM_RGM: Resonating Group Method for scattering	98,304 cores @ 8 hrs	Titan	X	X	

- Ab initio Methods (CC, GFMC, NCSM) →pushing the limits to calculate larger nuclei
- Density Functional Theory → reasonable time to solution to calculate the entire mass table

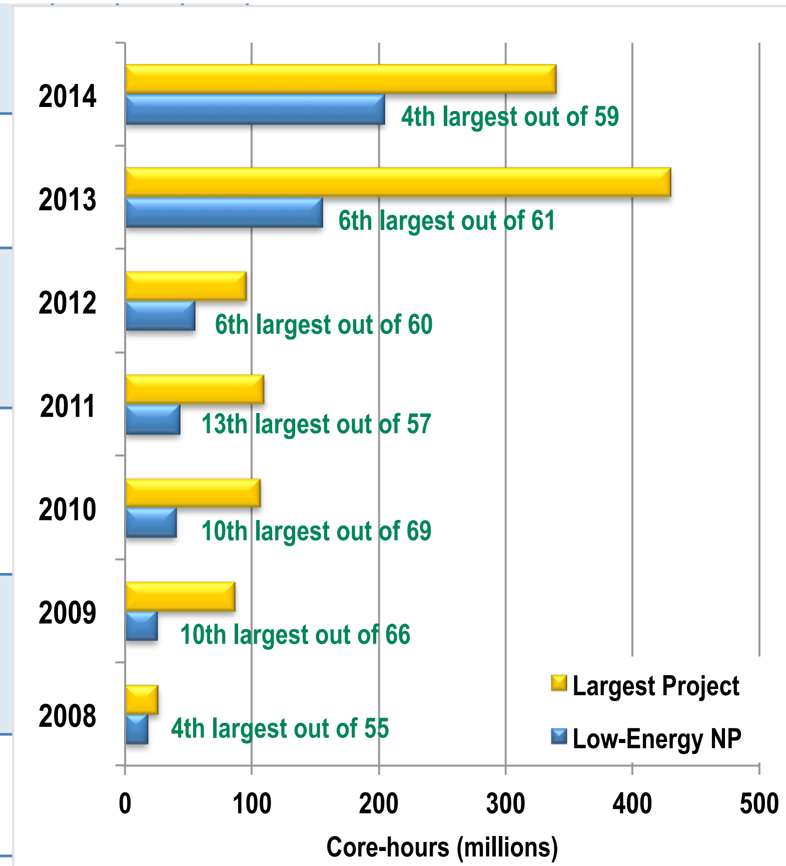

Slide by Hai Ah Nam, ORNL

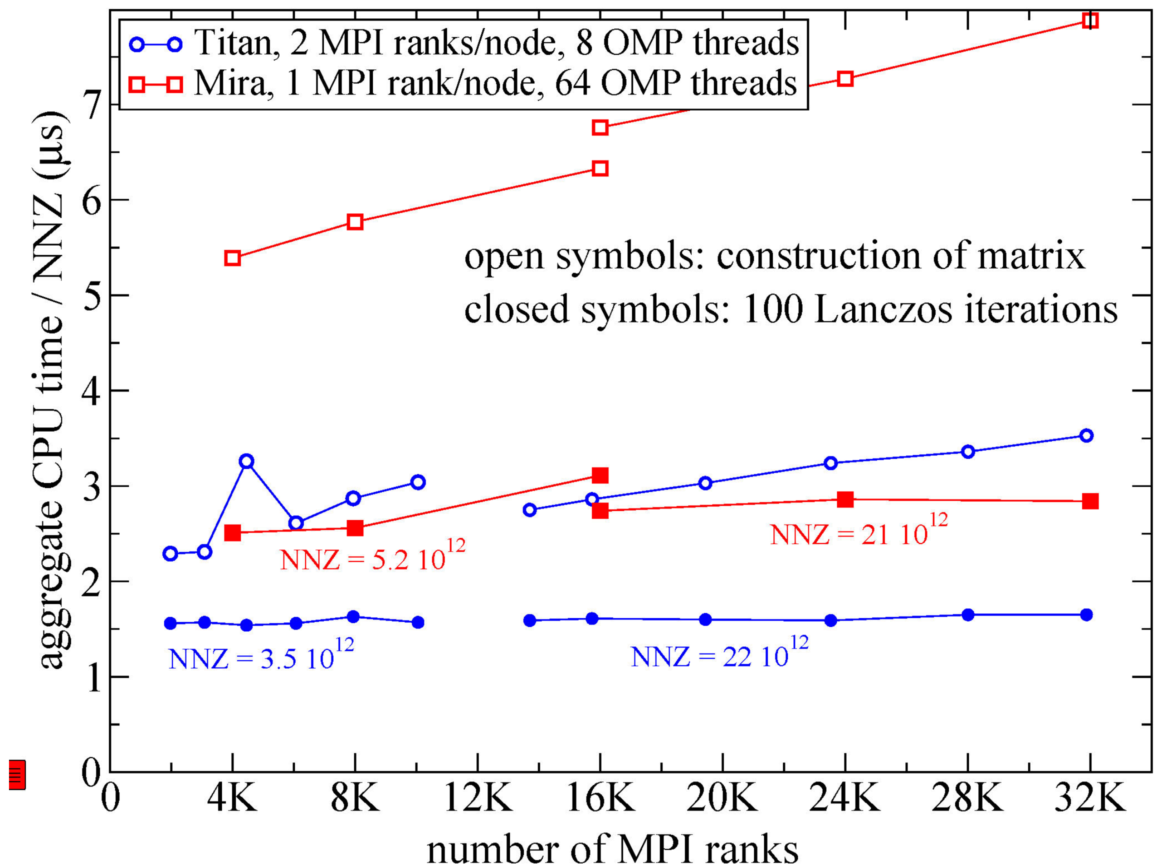

Strong and weak scaling for MFDn on Titan (without GPUs) and on MIRA. Aggregate CPU time per number of nonzero (NNZ) many-body matrix elements versus number of MPI ranks is shown for three test cases labeled by their NNZ values. All cases use 3N forces. One anomaly appears with the spike at 4K MPI ranks.

Note however - what worked well for light nuclei requires development/testing for heavier nuclei.

Leveraging GPUs in Ab Initio Nuclear Physics Calculations

Dossay Oryspayev*, Hugh Potter†, Pieter Maris†, Masha Sosonkina*‡, James P. Vary†, Sven Binder§, Angelo Calci§, Joachim Langhammer§, and Robert Roth§

IEEE 27th Parallel and Distributed Processing Symposium Workshops\& PhD Forum (IPDPSW), 1365 (2013)

Decouple NNN interaction matrix elements from JT-scheme to m-scheme

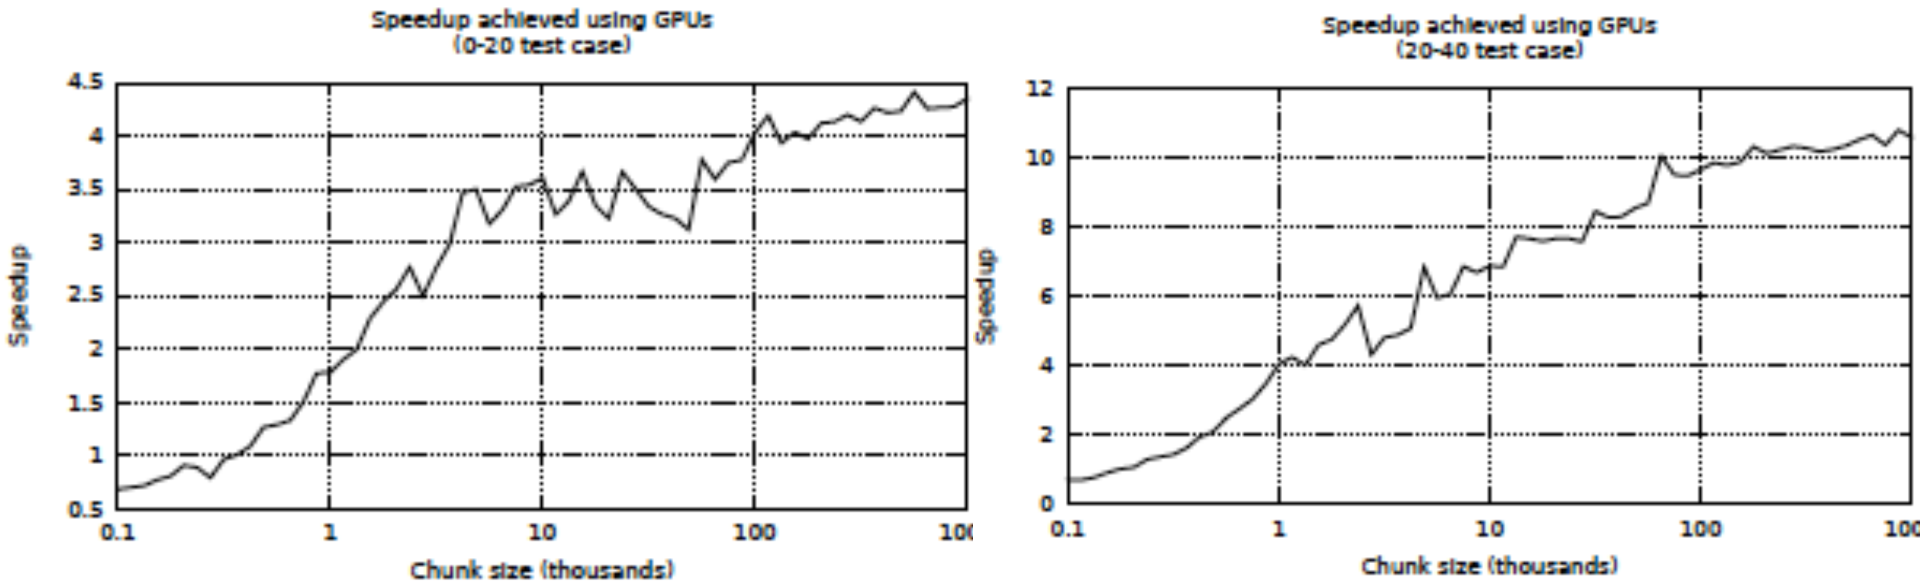

(a) 0-20 test case

(b) 20-40 test case

The bigger the workload transferred to the GPU, the greater the gain up to a limit

Scalable Eigensolver for MFDn

ASCR/NP – Applied Math/Computer Science Highlight

Objective

- Efficient and scalable iterative solvers for extreme-scale eigenvalue problems arising in nuclear physics

Impact

- Drastically reduced communication overheads
- Significant speed-ups over earlier version of MFDn (up to 6x on 18,000 cores)
- Almost perfect strong scaling on up to 260,000 cores on Jaguar

MFDn's comm. graph

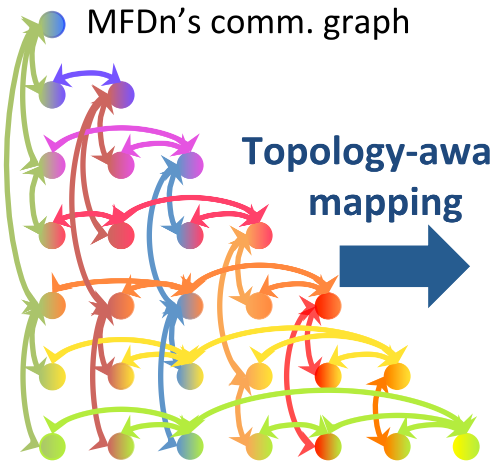

Topology-aware mapping

3D torus

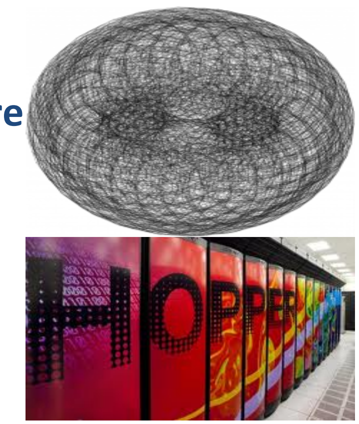

Topology-aware mapping of processes to the physical processors becomes more important as the gap between computational power and bandwidth widens. Communication groups are optimized through a column-major ordering of processes on the triangular grid [1].

Communication Hiding

Flow-chart for multi-threaded SpMV computations during the eigensolve phase of MFDn. Expensive communications are overlapped with computations. Explicit communications are carried out over topology-optimized groups [2].

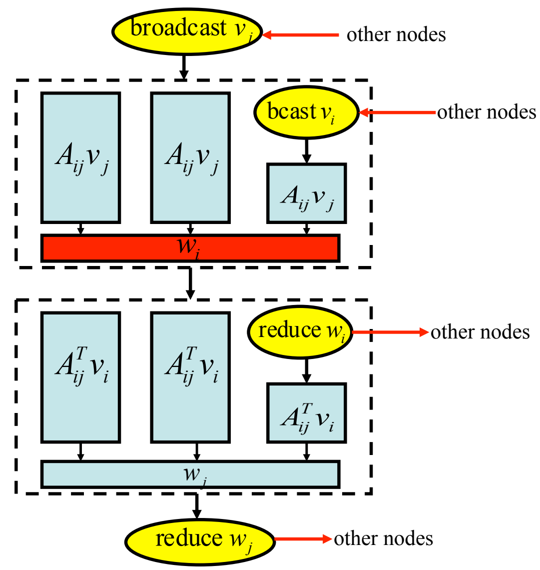

[1] H.M. Aktulga,C. Yang,P. Maris,J.P. Vary,E.G. Ng, "Topology-Aware Mappings for Large-Scale Eigenvalue Problems", Euro-Par 2012 Conference

[2] H.M. Aktulga,C. Yang,E.G. Ng,P. Maris,J.P. Vary, "Improving the Scalability of a Symmetric Iterative Eigensolver for Multi-core Platforms", CCP&E, in review

U.S. DEPARTMENT OF ENERGY | Office of Science

INCITE Advancing America's Science and Industrial Competitiveness

UNEDF SciDAC Collaboration Universal Nuclear Energy Density Functional

NUCLEI Nuclear Computational Low-Energy Initiative

Computational challenges of the NCSM – CPU time and Memory to store H

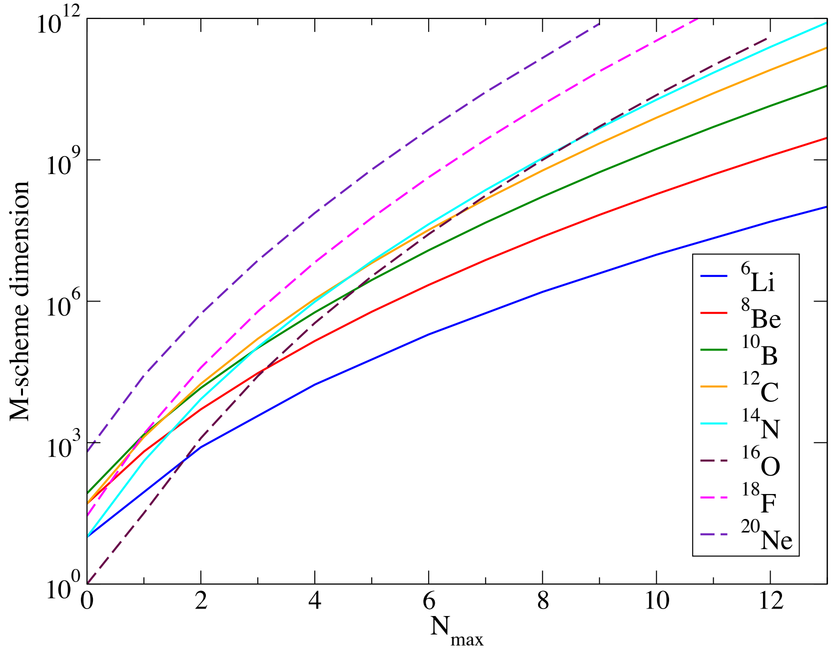

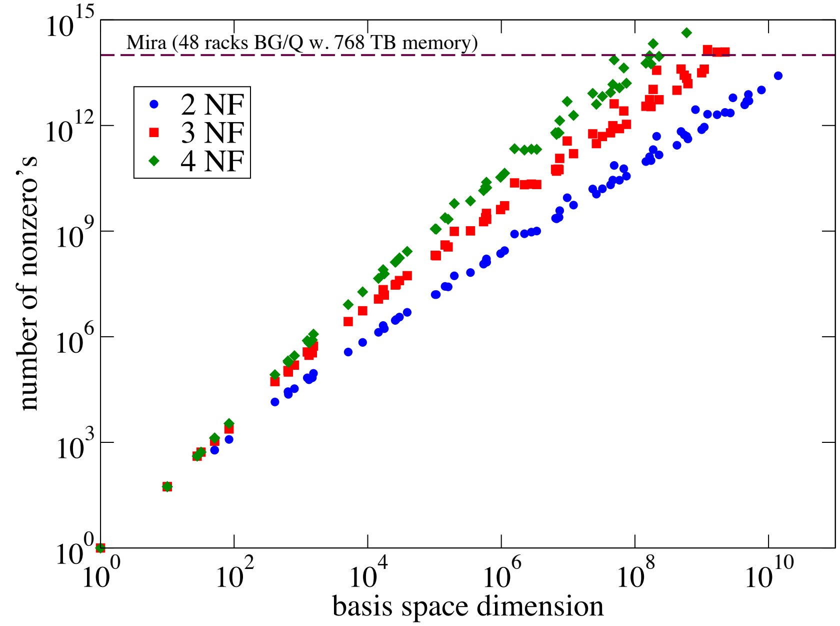

- Increase of basis space dimension with increasing A and N_{max}
  - need calculations up to at least $N_{\text{max}} = 8$
    for meaningful extrapolation and numerical error estimates
- More relevant measure for computational needs
  - number of nonzero matrix elements
  - current limit 10^{13} to 10^{14} (Hopper, Edison, Jaguar/Titan, Mira)

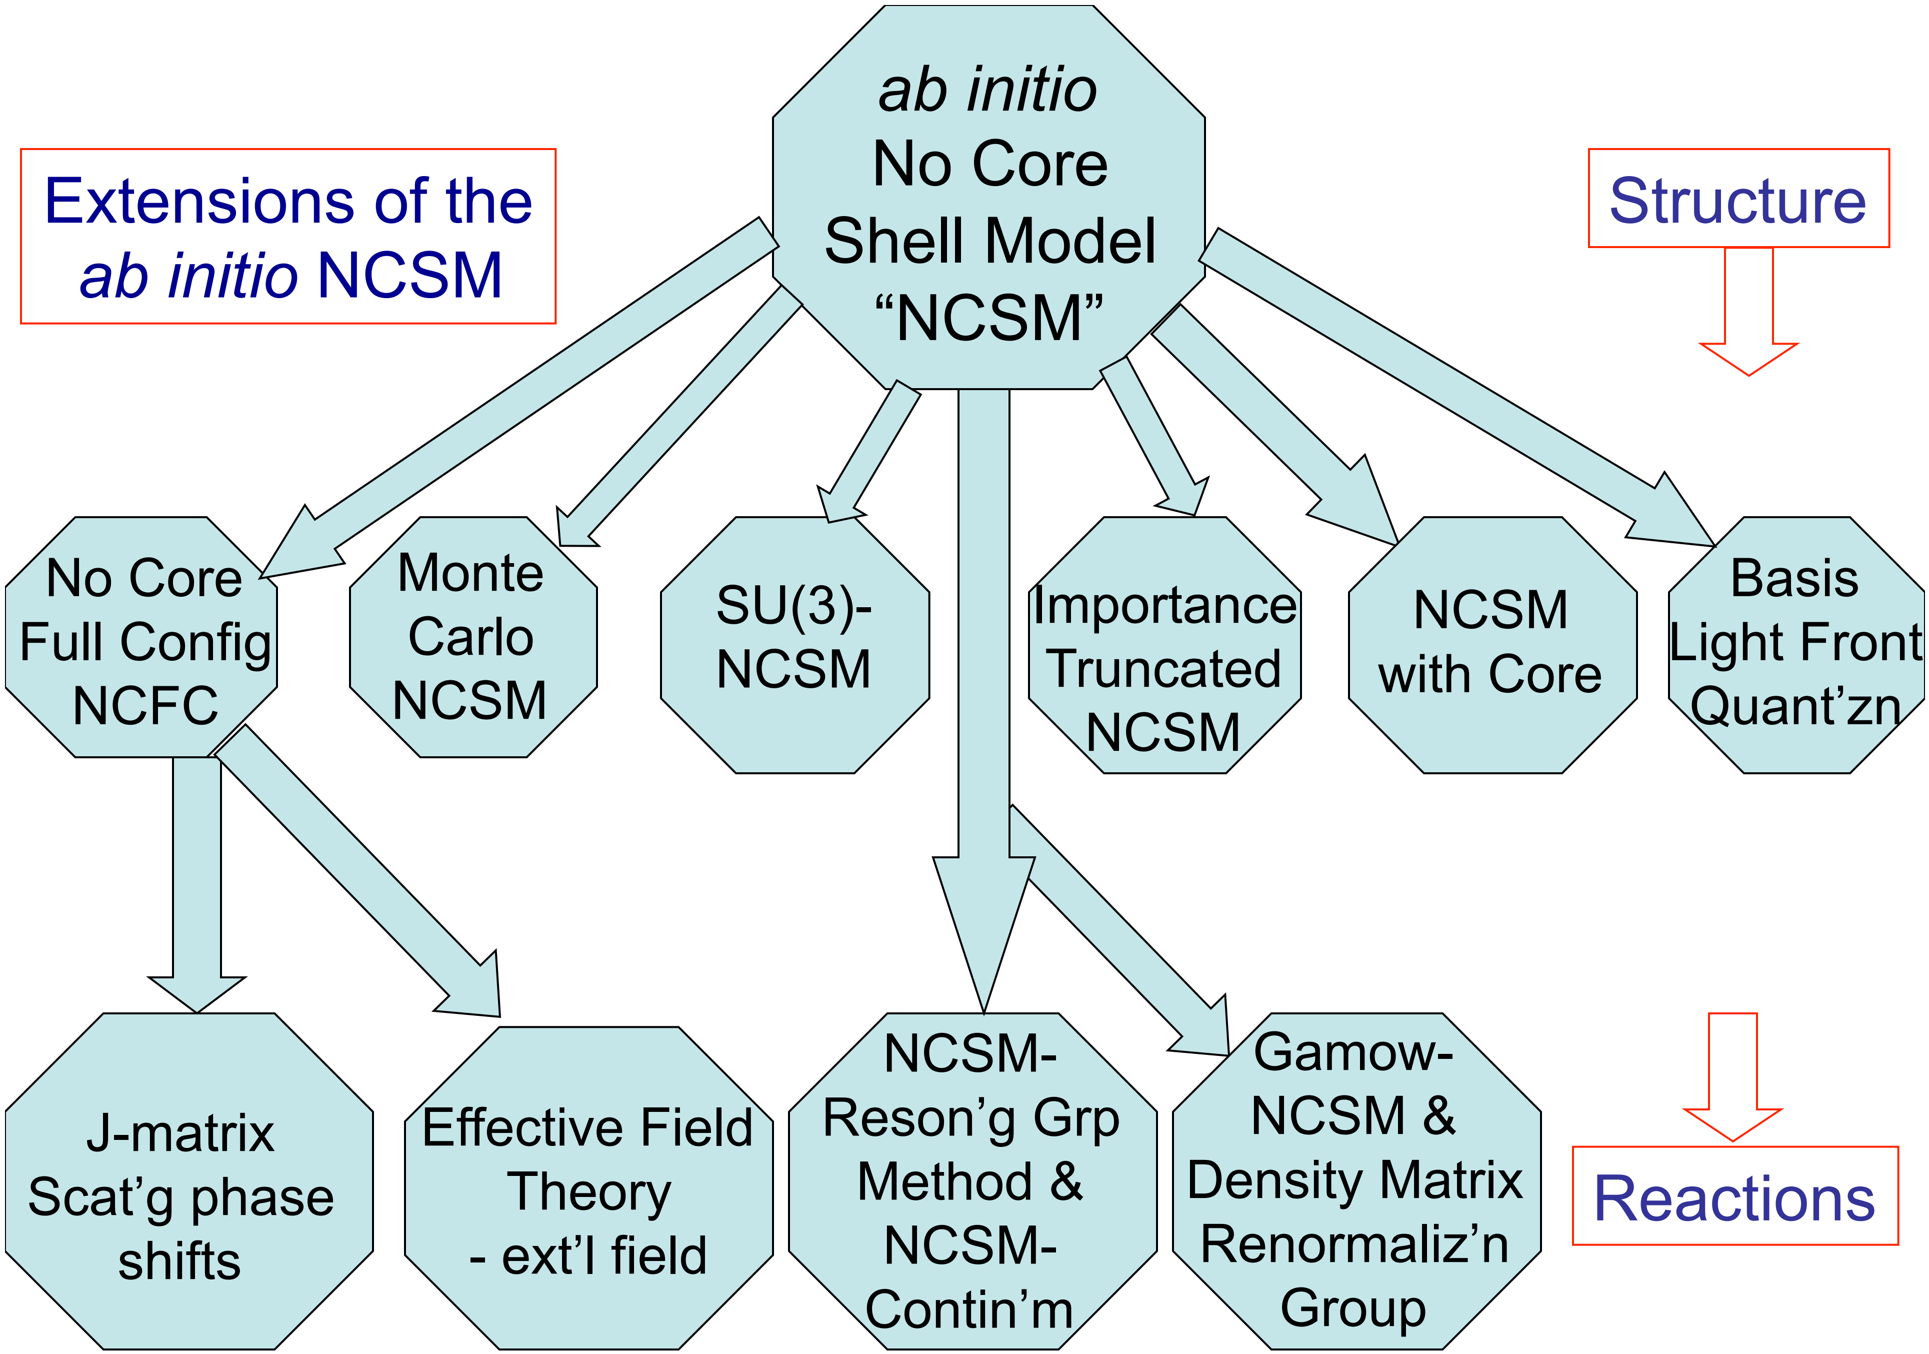
Extensions of the
ab initio NCSM
ab initio
No Core
Shell Model
“NCSM”
Structure
No Core
Full Config
NCFC
Monte
Carlo
NCSM
SU(3)-
NCSM
Importance
Truncated
NCSM
NCSM
with Core
Basis
Light Front
Quant’zn
J-matrix
Scat’g phase
shifts
Effective Field
Theory
- ext’l field
NCSM-
Reson’g Grp
Method &
NCSM-
Contin’m
Gamow-
NCSM &
Density Matrix
Renormaliz’n
Group
Reactions

spectrum A=8 nuclei with N3LO 2-body + N2LO 3-body

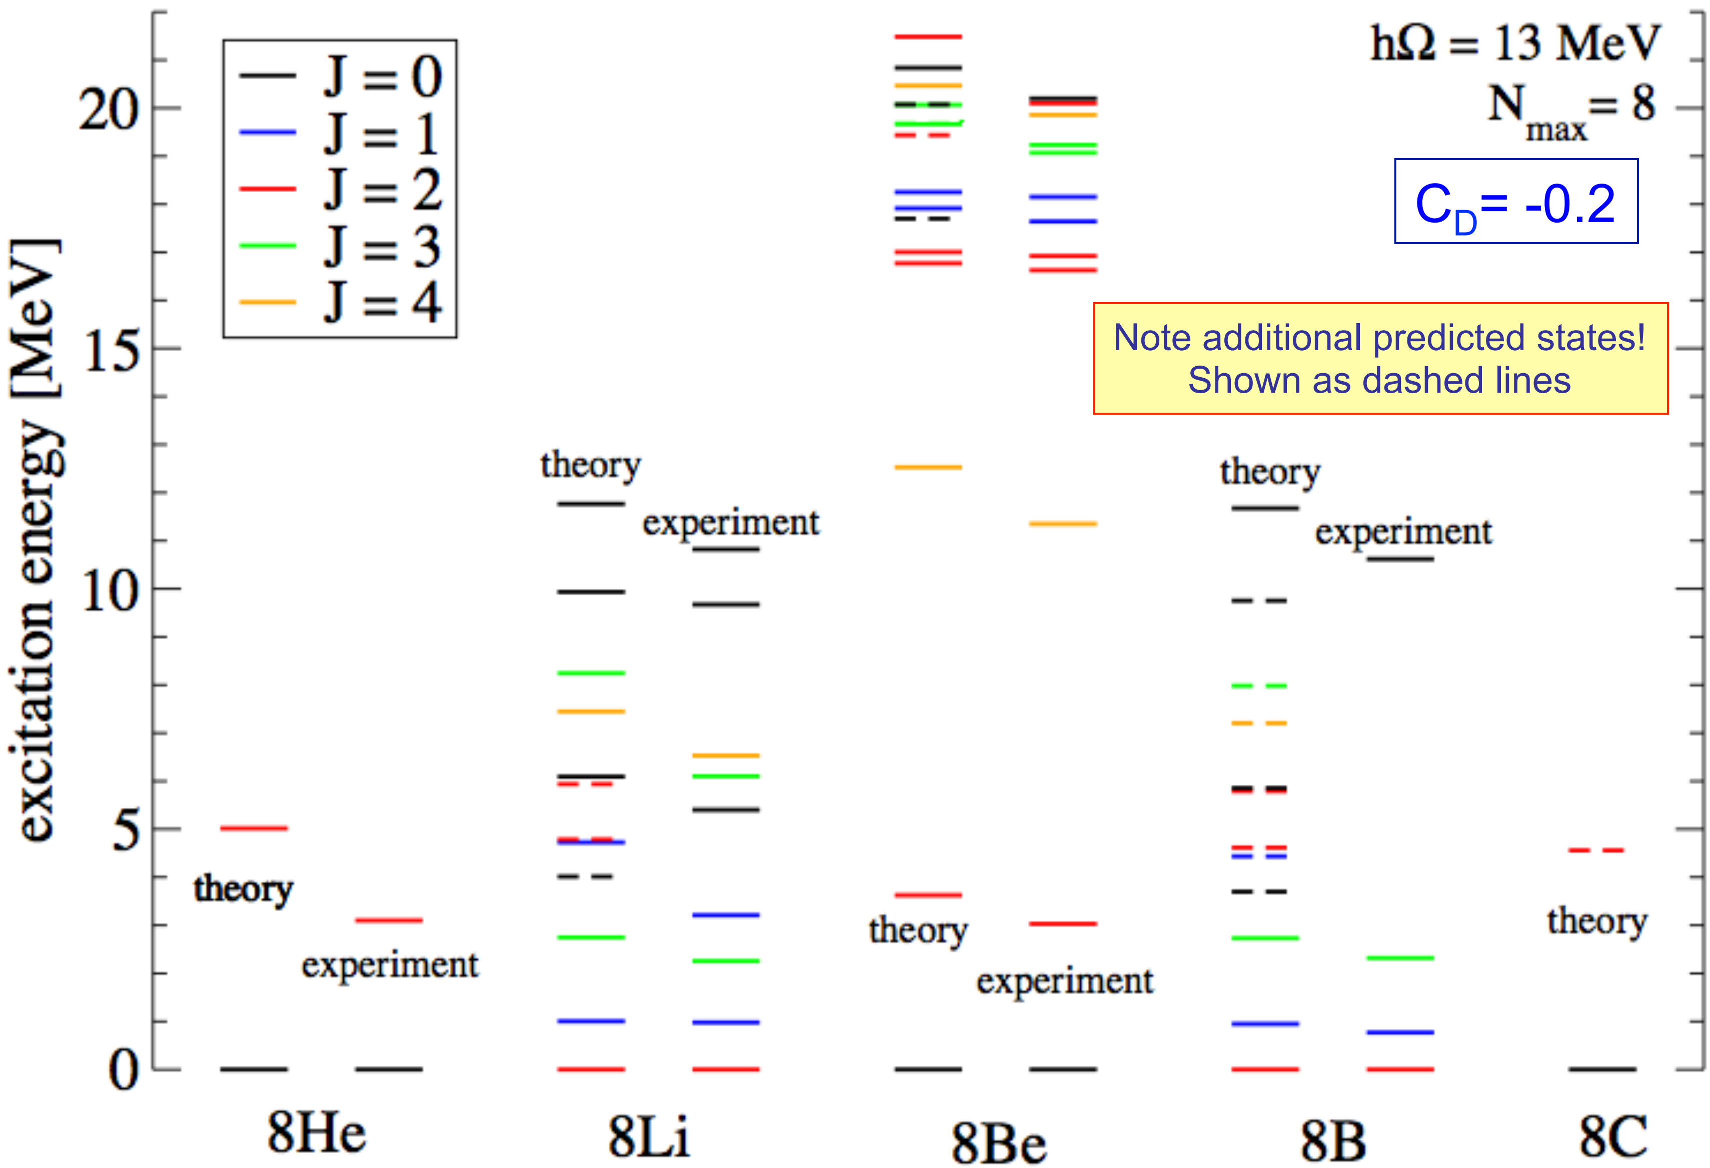

P. Maris, J. P. Vary and P. Navratil, Phys. Rev. C87, 014327 (2013); arXiv 1205.5686

9Be Translationally invariant gs density
Full 3D densities = rotate around the vertical axis

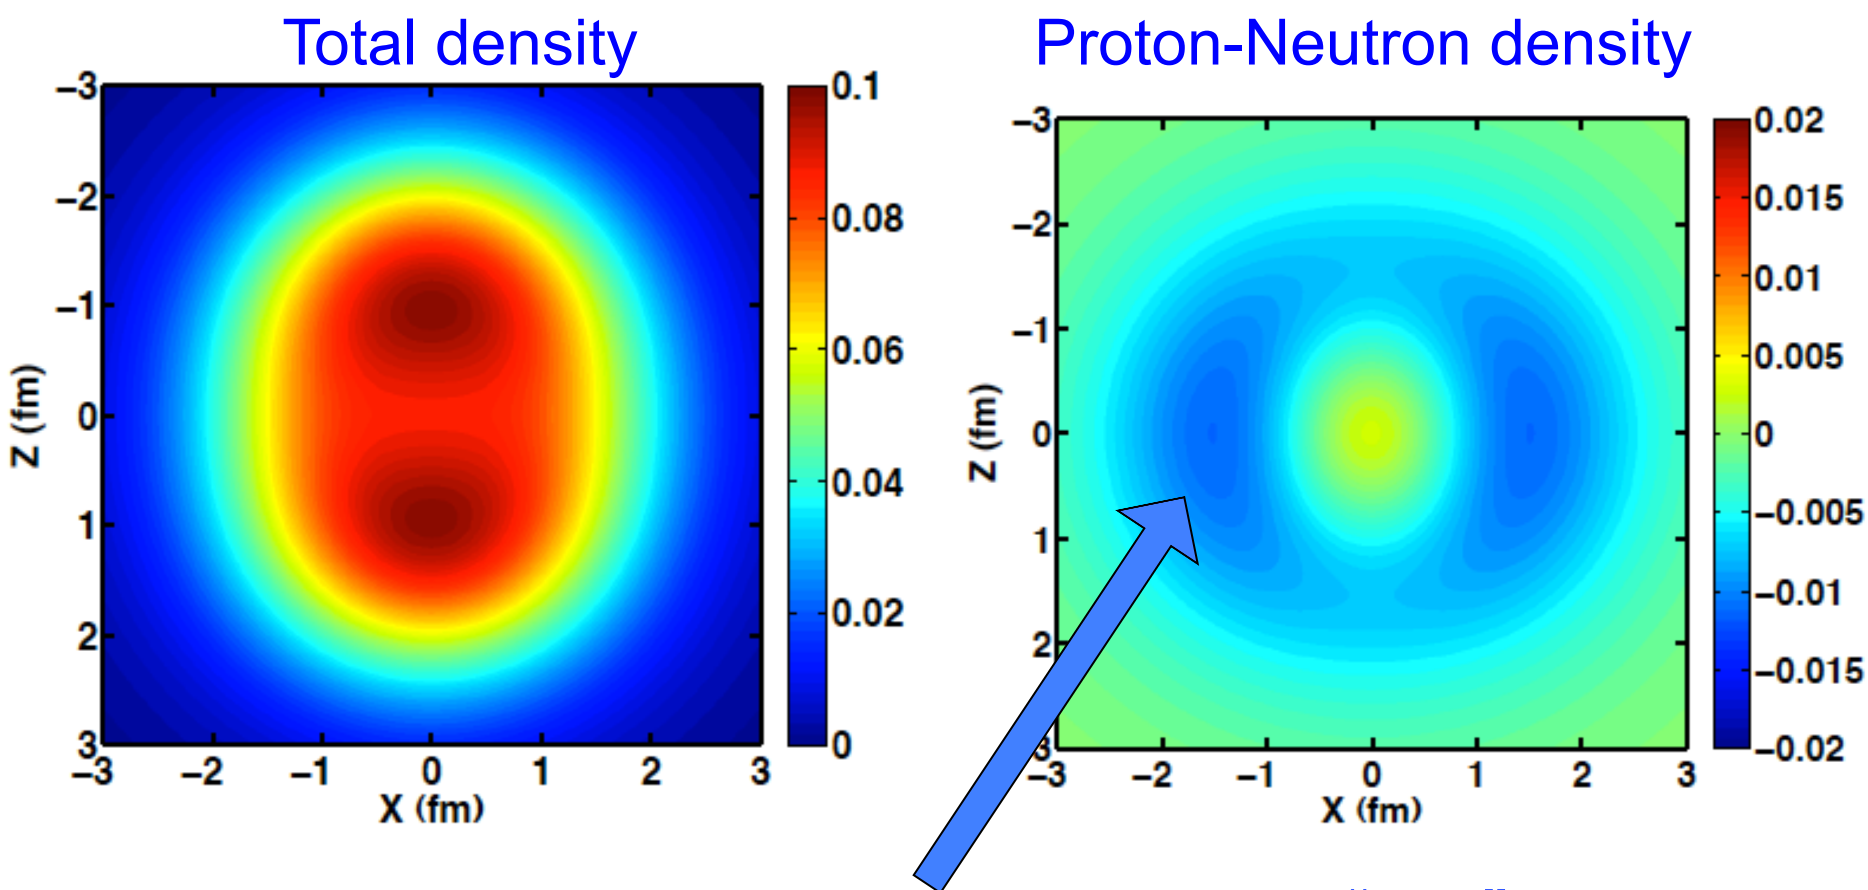

Shows that one neutron provides a “ring” cloud around two alpha clusters binding them together

C. Cockrell, J.P. Vary, P. Maris, Phys. Rev. C 86, 034325 (2012); arXiv:1201.0724;
C. Cockrell, PhD, Iowa State University

Ground state magnetic moments with JISP16

P. Maris and J.P. Vary, Int. J. Mod. Phys. E 22, 1330016 (2013)

$$\mu = \frac{1}{J+1}\Big(\langle \mathbf{J}\cdot\mathbf{L}_p\rangle + 5.586\langle \mathbf{J}\cdot\mathbf{S}_p\rangle - 3.826\langle \mathbf{J}\cdot\mathbf{S}_n\rangle\Big)\,\mu_0$$

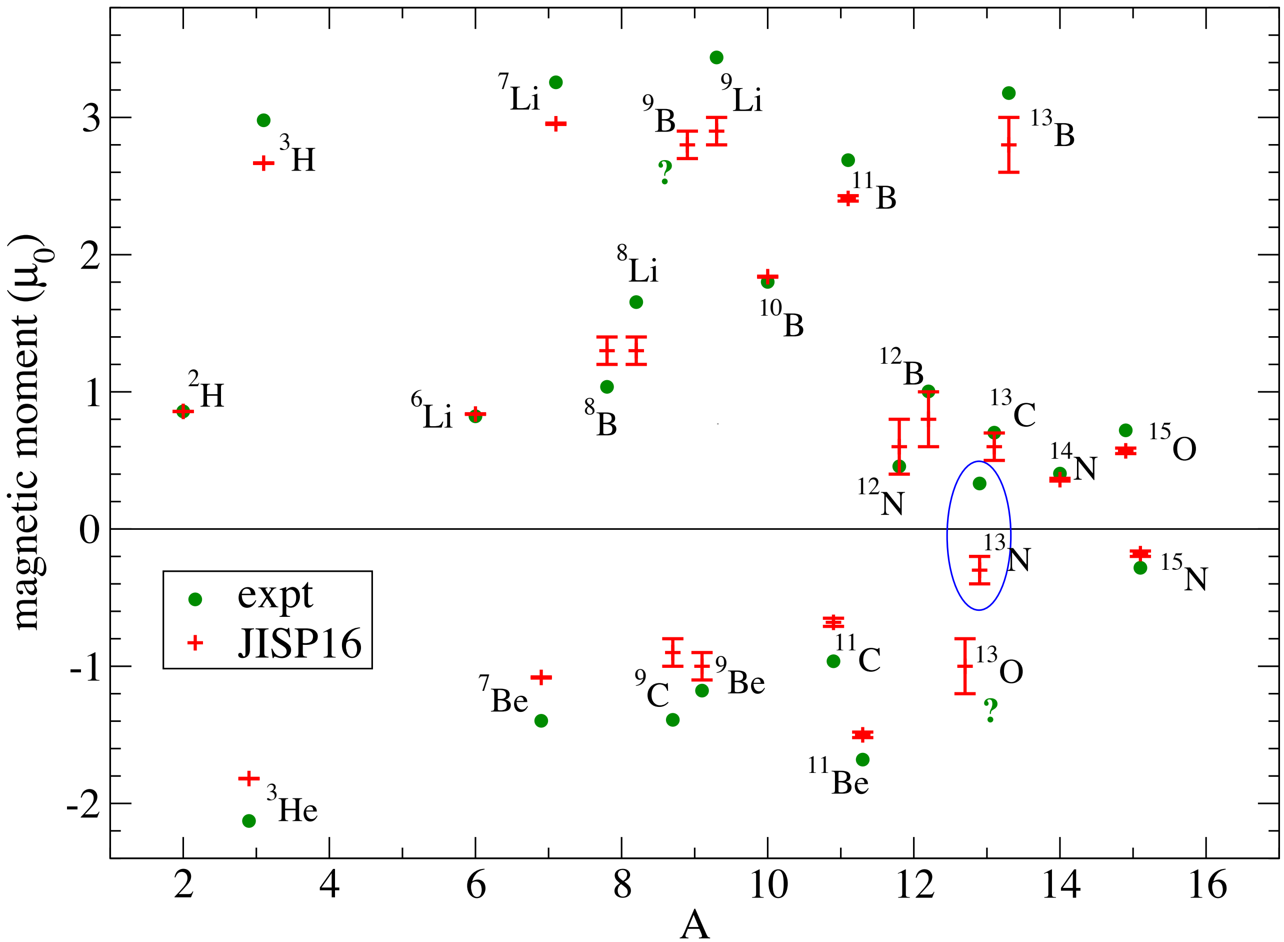

- Good agreement with data, given that we do not have any meson-exchange currents

How good is *ab initio* theory for predicting large scale collective motion?

Quantum rotator

$$E_J = \frac{\hat{J}^2}{2\mathbb{I}} = \frac{J(J+1)\hbar^2}{2\mathbb{I}}$$

$$\frac{E_4}{E_2} = \frac{20}{6} = 3.33$$

$$Experiment = 3.17$$

$$Theory(N_{\max} = 10) = 3.54$$

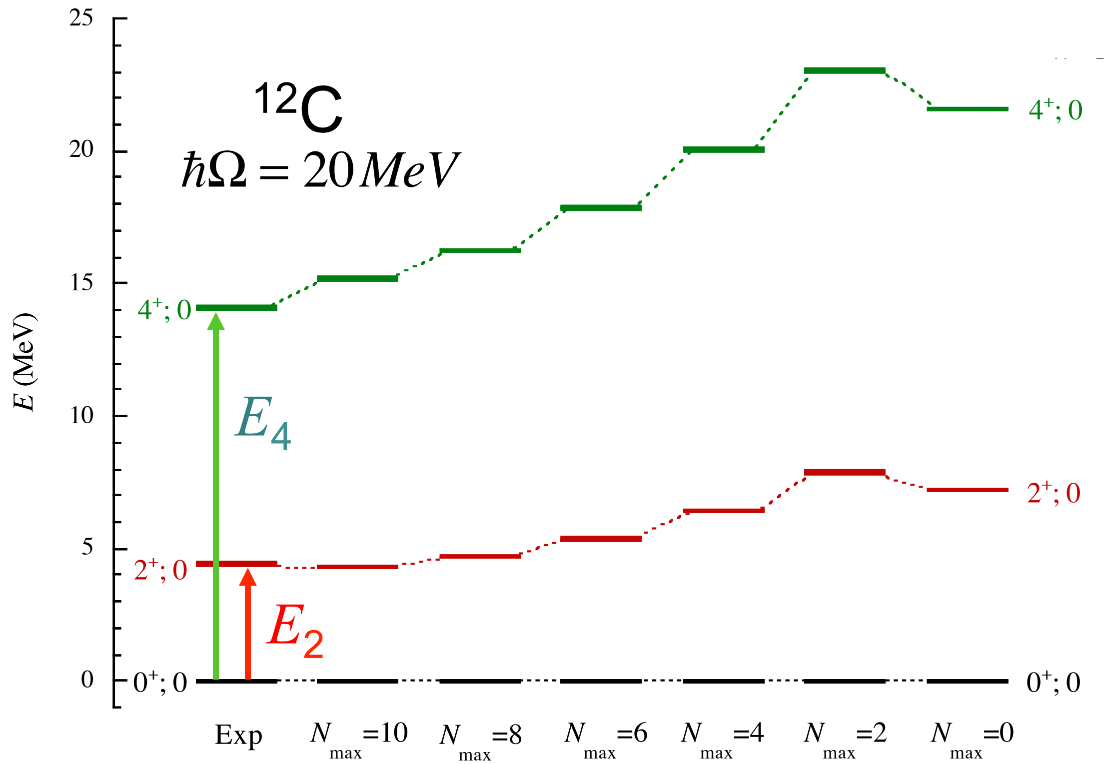

Dimension = 8×10^9

Extending the Reach of Ab Initio Applications:

Renormalization theory
Extrapolation theory
Physics-driven, theory-improved basis spaces
?

Optimize our utilization of available algorithms and computational resources
=> intense theoretical developments,
increase in the multi-disciplinary character, . . .

No-core shell model dimension

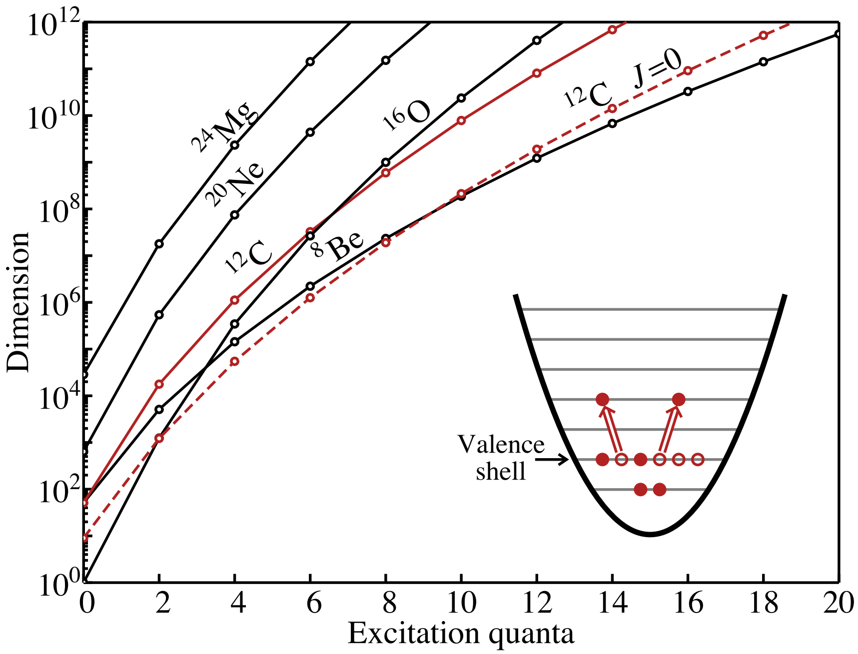

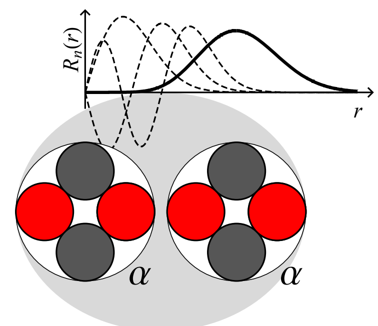

Cluster structure expected to require ~30–50$\hbar\Omega$ of oscillator excitation, *e.g.*, for $\alpha+\alpha+\alpha$ Hoyle state in ^{12}C.

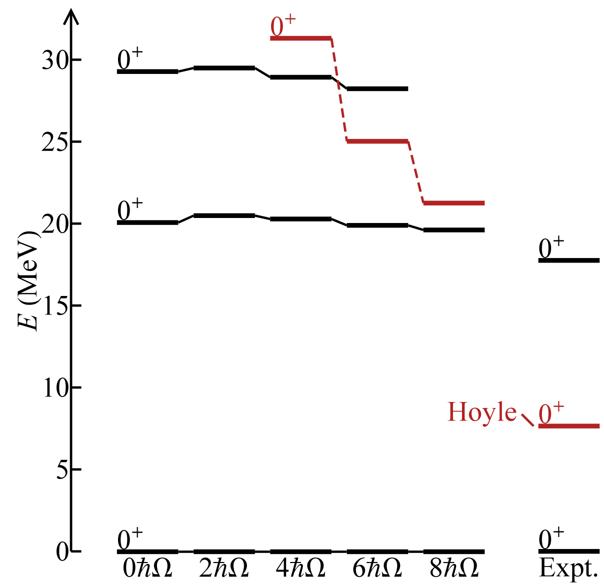

Calculations from T. Neff and H. Feldmeier, Eur. Phys. J. Special Topics **156**, 69 (2008).

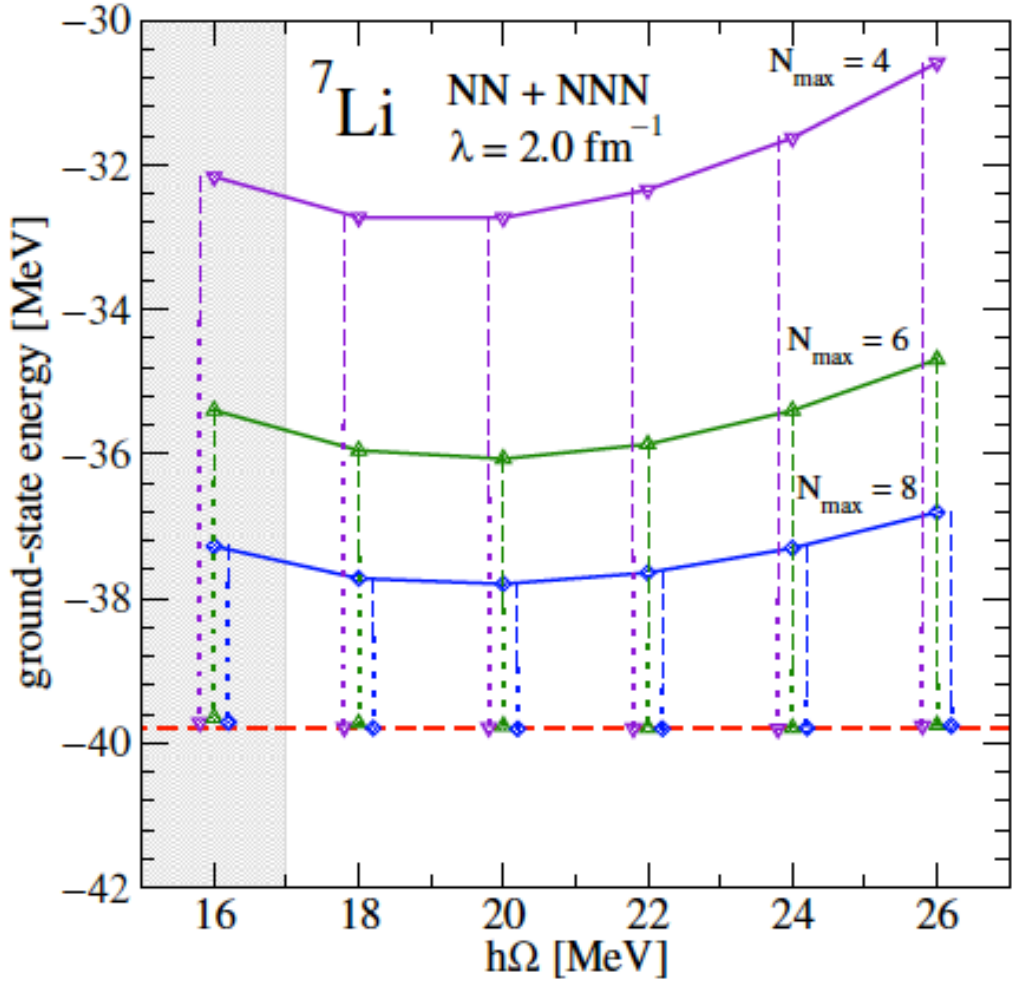

FIG. 17. (color online) Ground-state energy of ^{7}Li for the NN+NNN evolved Hamiltonians at $\lambda = 2.0\,\mathrm{fm}^{-1}$, with IR (vertical dashed) and UV (vertical dotted) corrections from Eq. (5) that add to predicted E_∞ values (points near the horizontal dashed line, which is the global E_∞).

E.D. Jurgenson, P. Maris, R.J. Furnstahl, P. Navratil, W.E. Ormand, J.P. Vary, Phys. Rev. C. 87, 054312 (2013); arXiv: 1302:5473

Ab initio Extreme Neutron Matter

Objectives

- Predict properties of neutron-rich systems which relate to exotic nuclei and nuclear astrophysics
- Determine how well high-precision phenomenological strong interactions compare with effective field theory based on QCD
- Produce accurate predictions with quantified uncertainties

Impact

- Improve nuclear energy density functionals used in extensive applications such as fission calculations
- Demonstrate the predictive power of ab initio nuclear theory for exotic nuclei with quantified uncertainties
- Guide future experiments at DOE-sponsored rare isotope production facilities

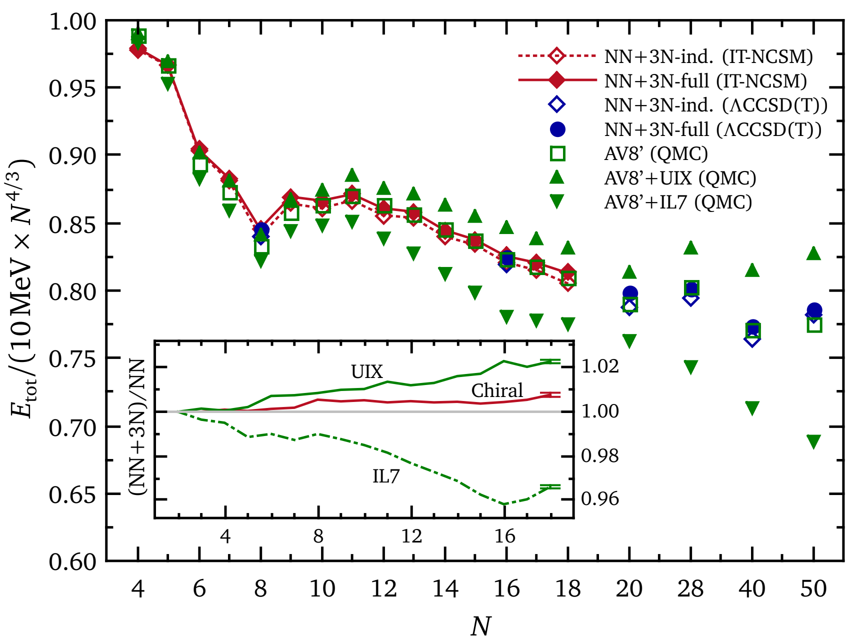

Comparison of ground state energies of systems with N neutrons trapped in a harmonic oscillator with strength 10 MeV. Solid red diamonds and blue dots signify new results with two-nucleon (NN) plus three-nucleon (3N) interactions derived from chiral effective field theory related to QCD. Inset displays the ratio of NN+3N to NN alone for the different interactions. Note that with increasing N, the chiral predictions lie between results from different high-precision phenomenological interactions, i.e. between AV8'+UIX and AV8'+IL7.

Accomplishments

1. Demonstrates predictive power of *ab initio* nuclear structure theory.
2. Provides results for next generation nuclear energy density functionals
3. Leads to improved predictions for astrophysical reactions
4. Demonstrates that the role of three-nucleon (3N) interactions in extreme neutron systems is significantly weaker than predicted from high-precision phenomemological interactions

U.S. DEPARTMENT OF ENERGY | Office of Science

NUCLEI Nuclear Computational Low-Energy Initiative

References: P. Maris, J.P. Vary, S. Gandolfi, J. Carlson, S.C. Pieper, Phys. Rev. C87, 054318 (2013); H. Potter, S. Fischer, P. Maris, J.P. Vary, S. Binder, J. Langhammer and R.Roth, in preparation; Contact: jvary@iastate.edu

Next Generation Ab Initio Applications:

Electroweak processes
Beyond the Standard Model
Neutrinoful and neutrinoless double beta-decay
?

Each puts major demands on theory, algorithms and computational resources
Growing demands => larger collaborating teams, growing computational resources,
Increase in the multi-disciplinary character, . . .

Origin of the anomalously long life-time of ^{14}C

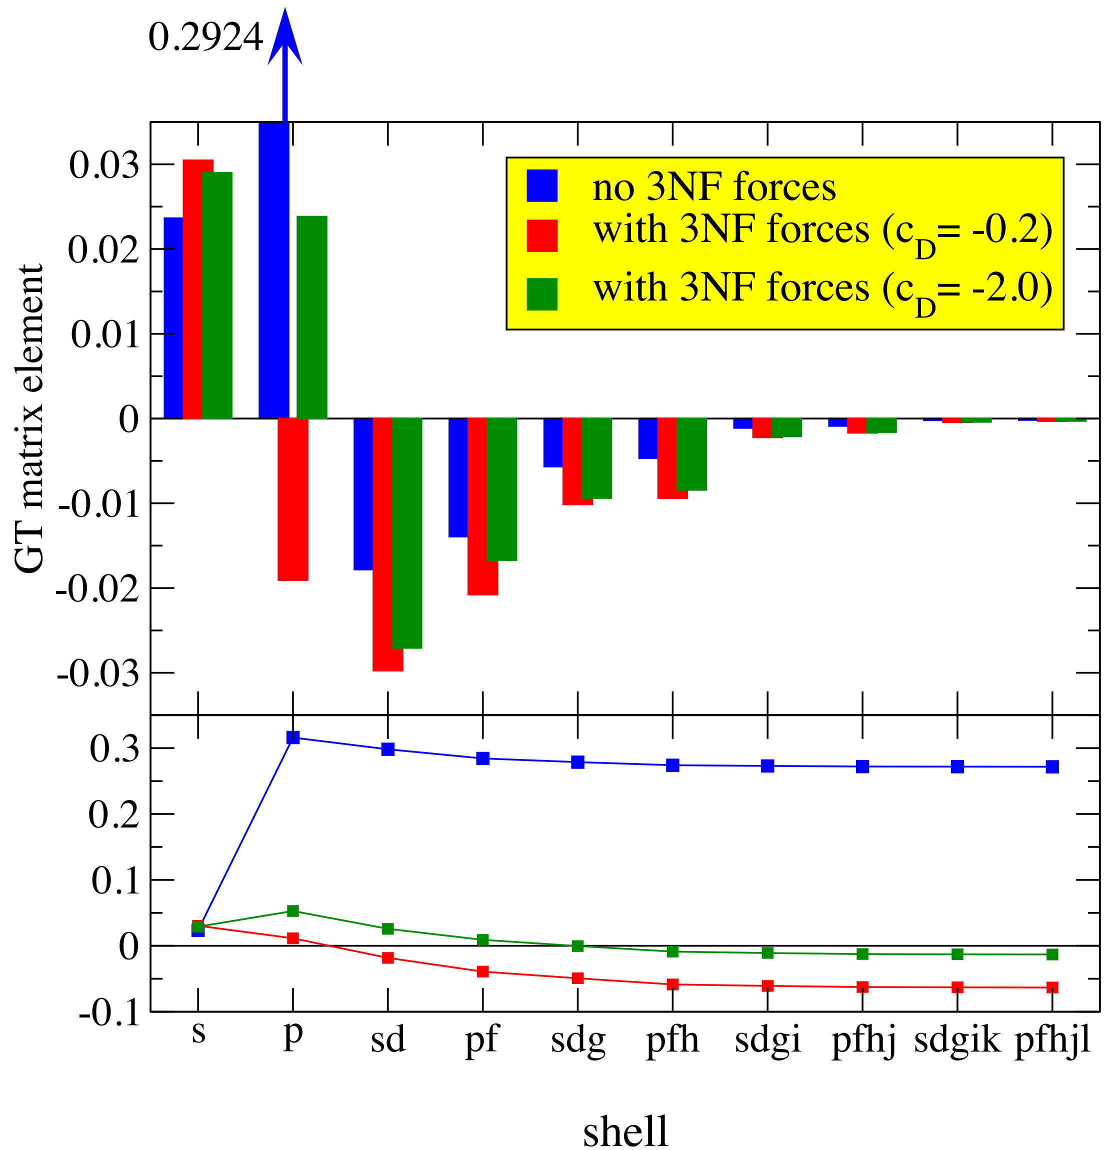

- near-complete cancellations between dominant contributions within p-shell
- very sensitive to details

Maris, Vary, Navratil, Ormand, Nam, Dean, PRL106, 202502 (2011)

Next Generation Systems and Resources

Moore's Law – Hardware
Moore's Law - Algorithms
Nuclear Theory's access to Systems and Computer Scientist Collaborators
?

Major demands on theory, algorithms and computational resources
Growing demands => larger collaborating teams, growing computational resources,
Increase in the multi-disciplinary character, . . .

Test of convergence: Use Hamiltonian QED in Basis Light Front Quantization to calculate the electron anomalous magnetic moment in e + e-gamma sectors

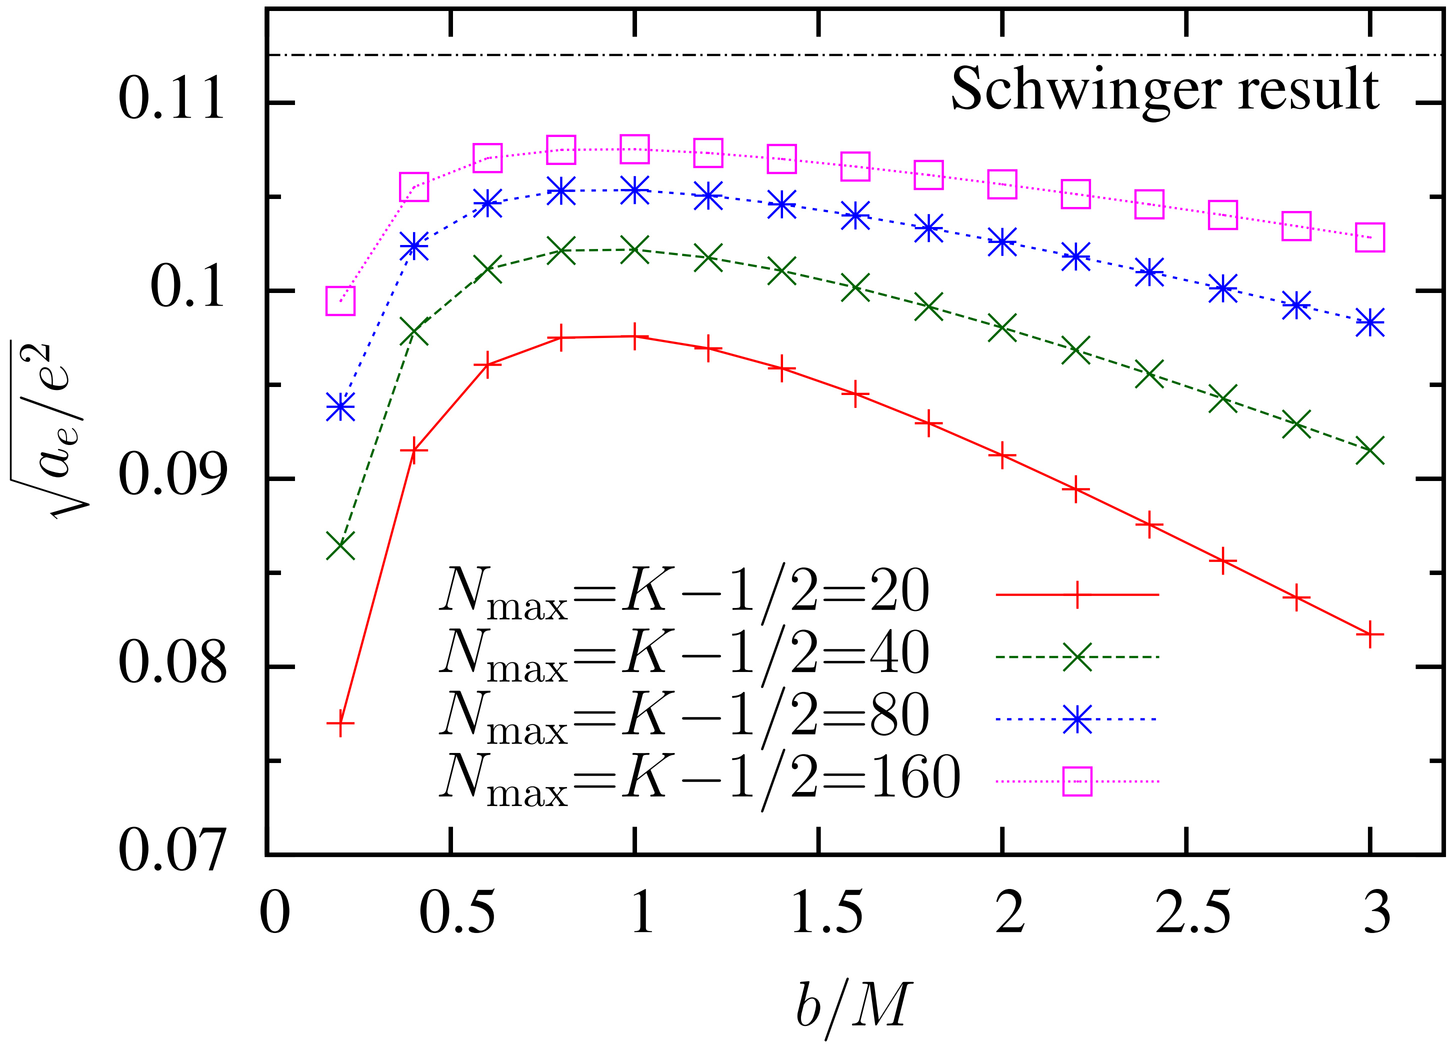

X. Zhao, H. Honkanen, P. Maris, J. P. Vary and S. J. Brodsky , arXiv 1402.4195

Test of convergence: Use Hamiltonian QED in Basis Light Front Quantization to calculate the electron anomalous magnetic moment in e + e-gamma sectors

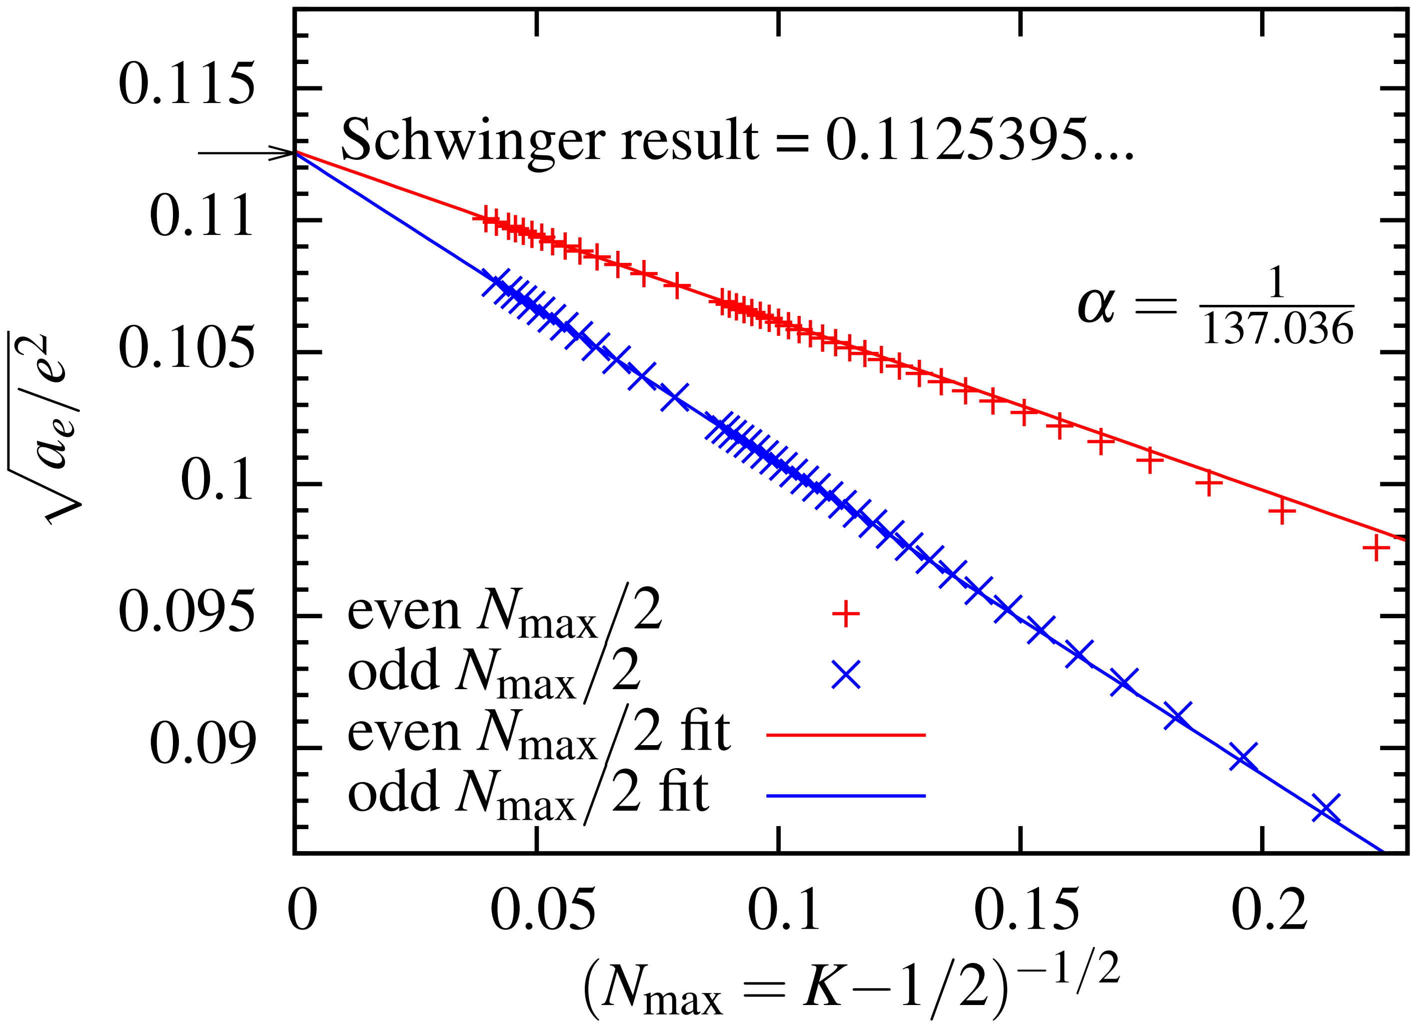

X. Zhao, H. Honkanen, P. Maris, J. P. Vary and S. J. Brodsky , arXiv 1402.4195

Many recent insights obtained from ab initio NCSM/NCFC:

Collective modes in light nuclei accessible with ab initio approach
3NFs continue to play an important role in many observables
Neutron drop results show (sub)shell closures
IR and UV convergence in HO basis (Coon et al., Papenbrock et al.)
Alternative basis spaces poised to relieve IR shortcomings of HO basis
Alternative MB methods poised to access clustering, halo physics regions
Computer Science and Applied Math collaborations invaluable
Generous allocations of computer resources essential to progress

Many outstanding nuclear physics puzzles
and discovery opportunities

Clustering phenomena
Origin of the successful nuclear shell model
Nuclear reactions and breakup
Astrophysical r/p processes & drip lines
Predictive theory of fission
Existence/stability of superheavy nuclei
Physics beyond the Standard Model
Possible lepton number violation
Spin content of the proton
+ Many More!

6 Requirements Summary Worksheet

Please try to fill out this worksheet, based on your answers above, to be best of your ability prior to the review.

	Used at NERSC in 2013	Needed at NERSC in 2017
Computational Hours	27 (mostly Edison pre-acceptance usage)	72
Typical number of cores* used for production runs	From 5% to full machine	From 5% to full machine
Maximum number of cores* that can be used for production runs	Full machine	Full machine
Data read and written per run	< 1 TB	< 1 TB
Maximum I/O bandwidth	Not known	Not known
Percent of runtime for I/O	< 5%	< 5%
Scratch File System space	1 TB	3 TB
Shared filesystem space	2 TB	8 TB
Archival data	5 TB	50 TB
Memory per node	All available GB	Maximum possible GB
Aggregate memory	Full machine	Full machine