

Nuclear Structure and Reactions

Joe Carlson, Witek Nazarewicz, Rusty Lusk

- Introduction
- Recent Highlights
- Needs for the Future
- Ties: NP / ASCR / NNSA



SciDAC

Scientific Discovery through Advanced Computing



UNEDF SciDAC Collaboration
Universal Nuclear Energy Density Functional

Physics of Hadrons


Degrees of Freedom



quarks, gluons




constituent quarks




baryons, mesons

ab initio



protons, neutrons

CI



nucleonic densities and currents

DFT



collective coordinates

collective models

Energy (MeV)

940
neutron mass


140
pion mass

8
proton separation
energy in lead

1.12
vibrational
state in tin

0.043
rotational
state in uranium

Resolution




Effective Field Theory



Hot and dense quark-gluon matter

Hadron structure

Hadron-Nuclear interface



Nuclear structure
Nuclear reactions

Nuclear astrophysics
New standard model
Applications of nuclear science

Interfaces provide crucial clues

The Nuclear Landscape and the Big Questions (NAS report)

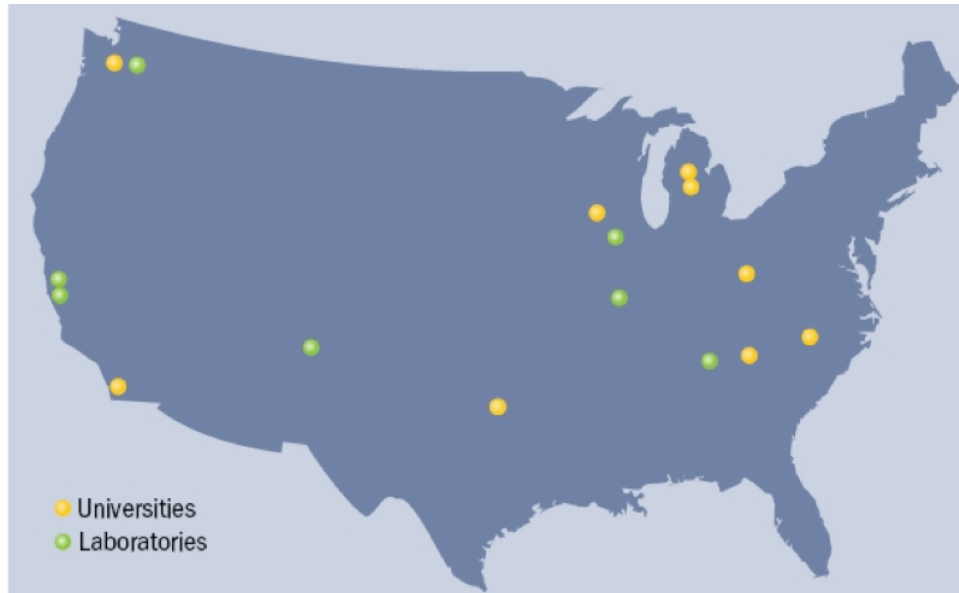
- How did visible matter come into being and how does it evolve?
- How does subatomic matter organize itself and what phenomena emerge?
- Are the fundamental interactions that are basic to the structure of matter fully understood?
- How can the knowledge and technological progress provided by nuclear physics best be used to benefit society?

Experimental relevance:

FRIB, LBNL Facilities,
NNSA facilities, JLab, JINA,
SNS, ...

EXAMPLE: *Universal Nuclear Energy Density Functional*

(other, smaller, collaborations exist: NuN, TORUS, PetaApps,...)



- Ab initio structure
- Ab initio functionals
- DFT applications
- DFT extensions
- Reactions

Focus on:

- Predictive power
- Robust extrapolations
- Validation
- Guidance

- Funded for 5 years by DOE (NP/SC, NNSA, ASCR)
- 9 universities and 7 national labs
- Junior scientists: 11 students, 19 postdocs/year
- ~50 researchers in
 - physics
 - computer science
 - applied mathematics
- International partners

For a popular description of UNEDF, see:

- SciDAC Review Winter 2007

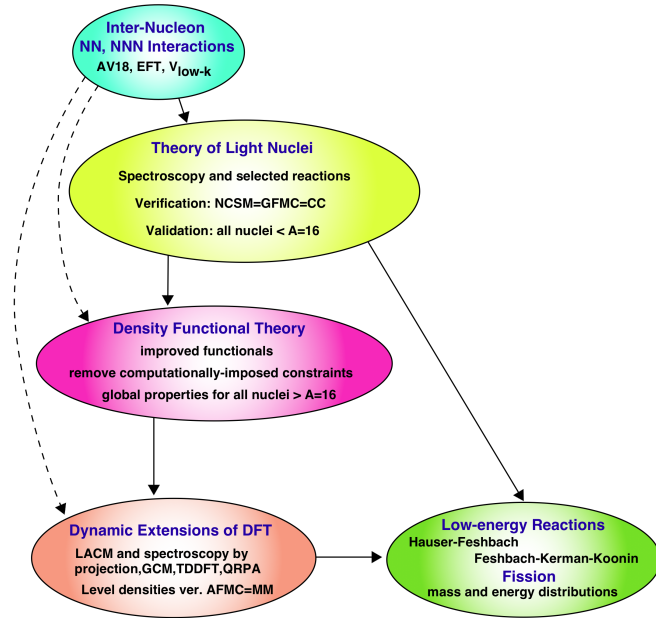
<http://www.scidacreview.org/0704/pdf/unedf.pdf>

- Nucl. Phys. News 21, No. 2, 24 (2011)

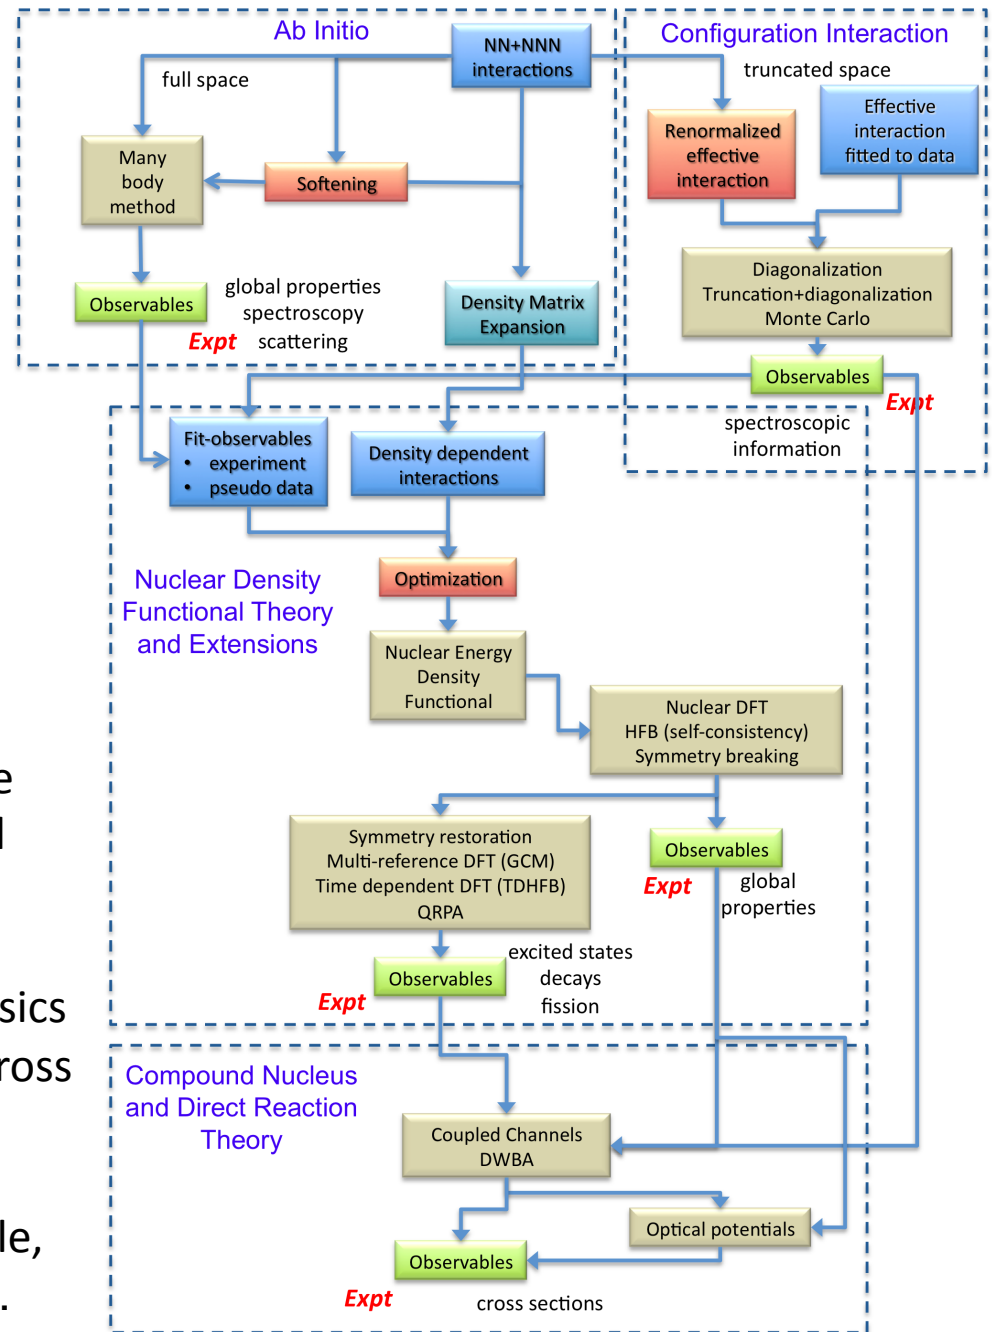
- Office of Science “Highlight Series”:

<http://science.energy.gov/news/in-focus/2011/03-28-11-s/>

UNEDF in 2007



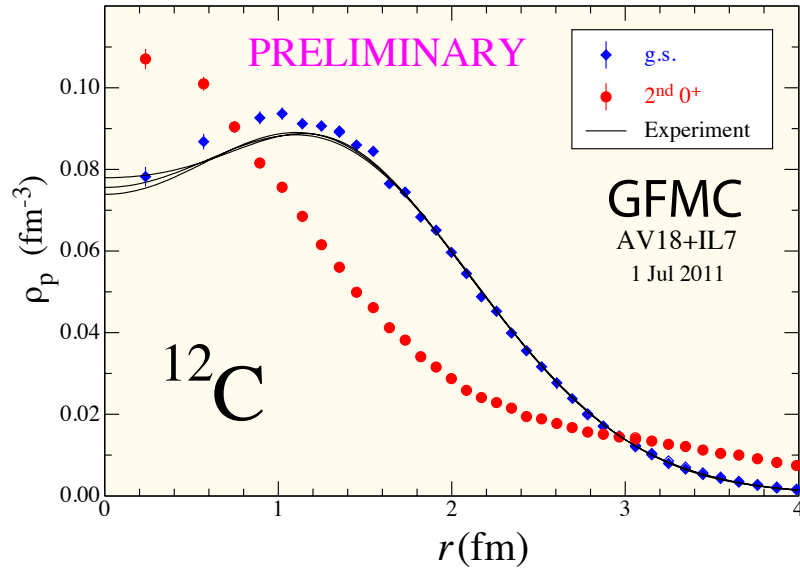
UNEDF in 2011




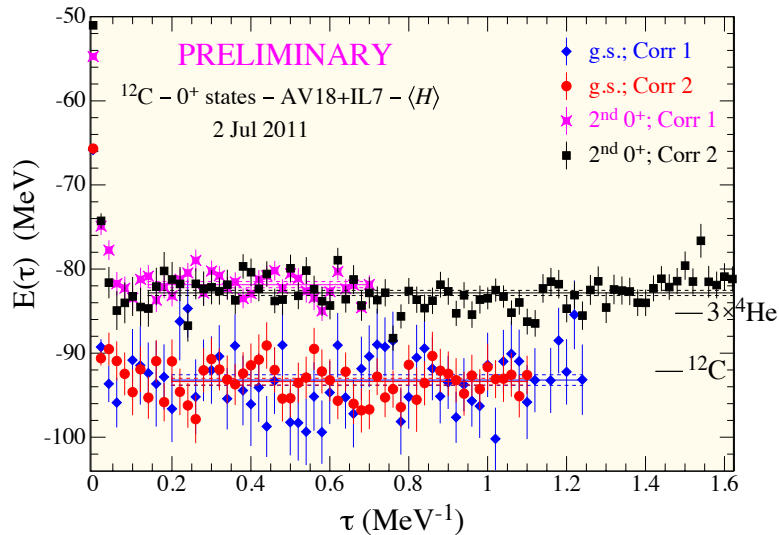
- As UNEDF has matured, it has been satisfying to see the increased coherence within the effort; the project has created and facilitated an increasing interplay among the major sections *where none existed previously*. Each of the main physics areas now includes collaborations that cross over into other sectors.
- A fair number of applications are now running on the biggest machines available, a vast change from when UNEDF started.

Ab initio description of ^{12}C


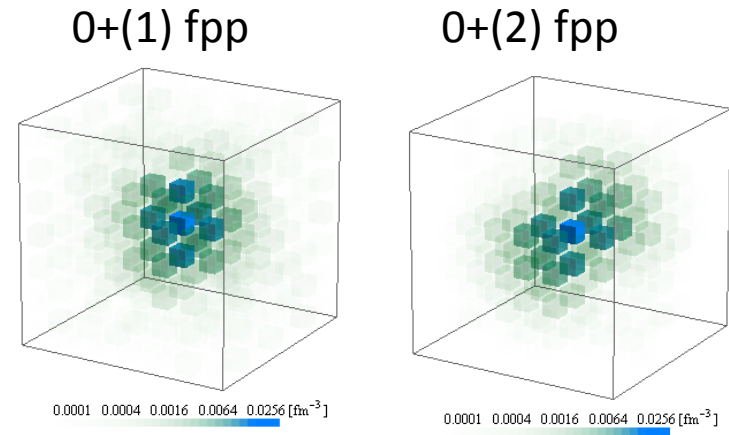
GFMC (Pieper et al.)



The ADLB (Asynchronous Dynamic Load-Balancing) library & GFMC. GFMC energy 93.5(6) MeV; expt. 92.16 MeV. GFMC pp radius 2.35 fm; expt. 2.33 fm.



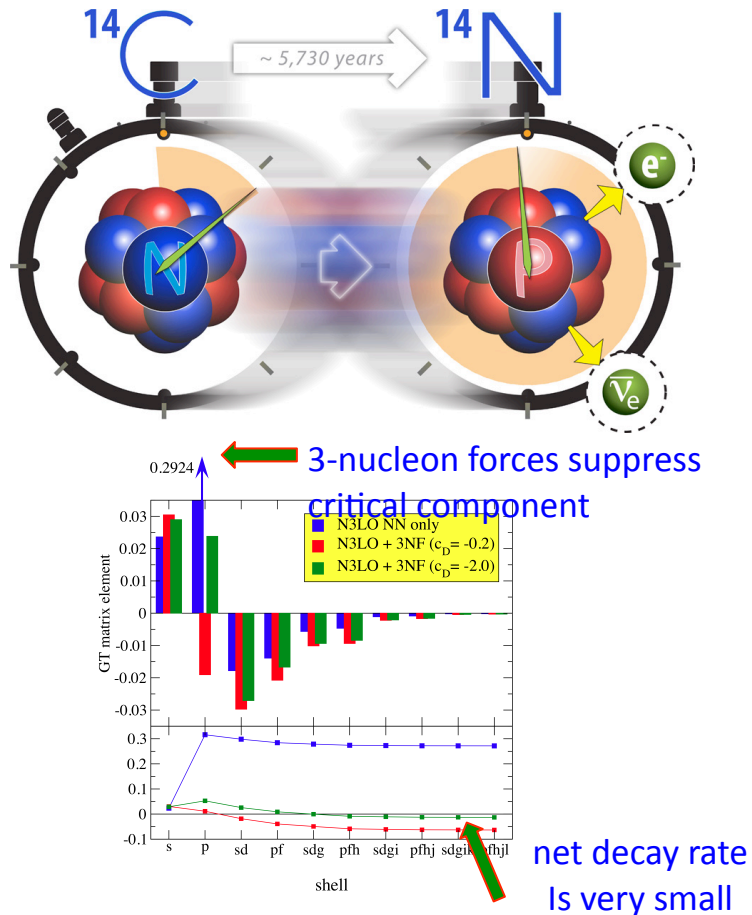
Lattice EFT (Lee, Epelbaum, Meissner,...)



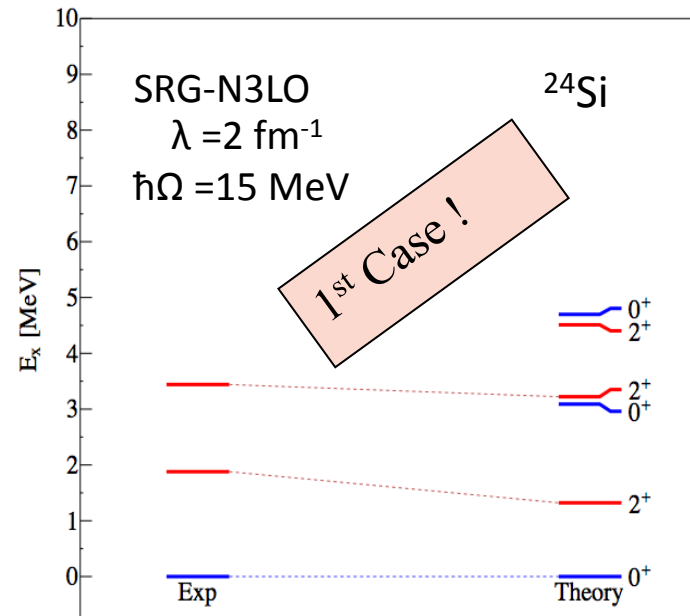
Epelbaum et al., Phys. Rev. Lett. 106, 192501 (2011)

Anomalous Long Lifetime of Carbon-14

Anomalous long lifetime of Carbon-14 (used in carbon dating) explained by ab-initio calculations using NN and NNN forces. *Phys. Rev. Lett.* 106, 202502 (2011)


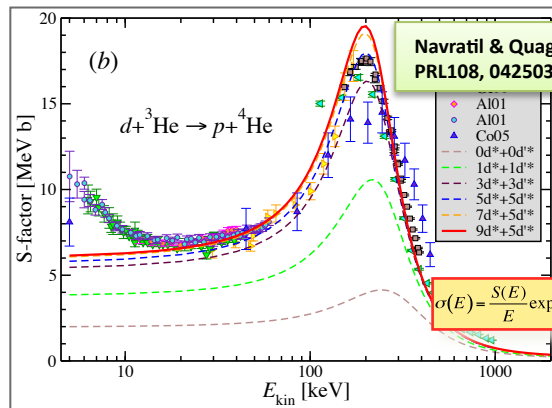


Truncating exploding Shell Model space:
Symmetry Adapted NCSM
Extending NCSM to Si



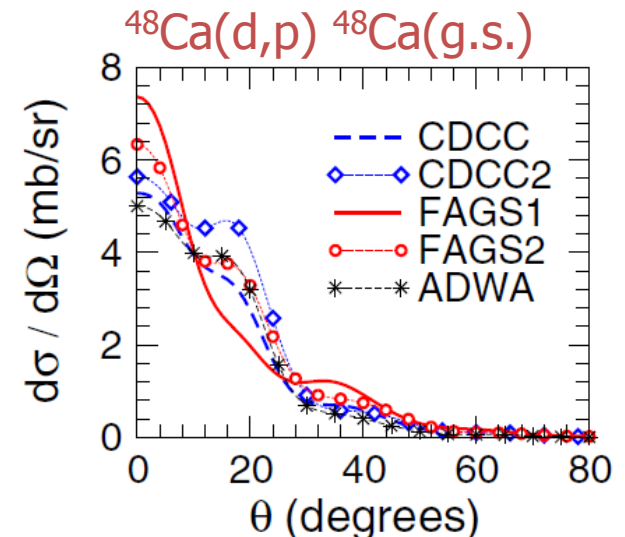
<http://www.newscientist.com/article/mg21128214.400-quantum-quirk-makes-carbon-dating-possible.html>
<http://phys.org/news/2011-05-physicists-lifetime-carbon-.html>

Nuclear reactions



NCSM/RGM is pioneering *ab initio* calculations of light-nuclei fusion reaction with NN interaction. Here, ${}^3\text{He}$ (d,p) ${}^4\text{He}$ S-factor.

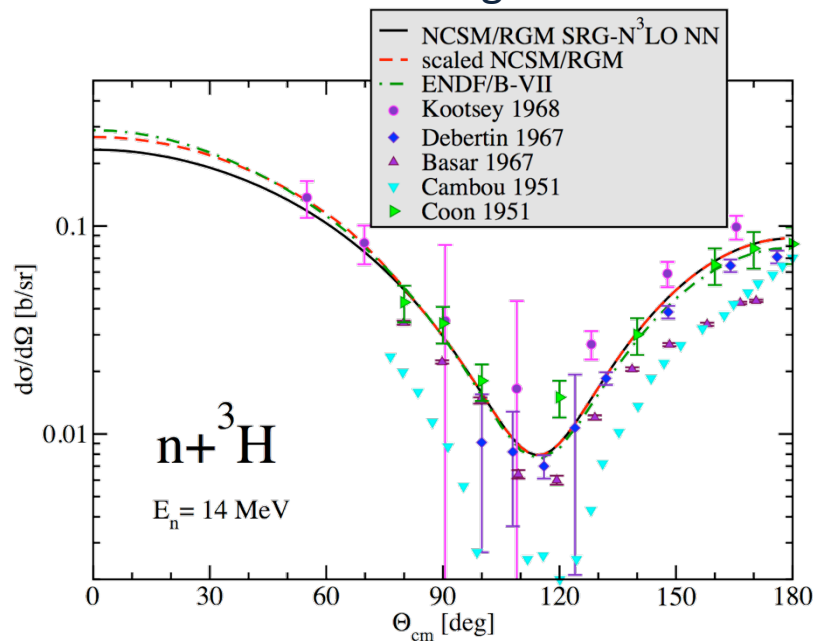
TORUS topical collaboration



PRC 84, 034607(2011), PRC 85, 054621 (2012)

Data deviate from NCSM/RGM results at low energy due to lab. electron-screening

Ab initio theory reduces uncertainty due to conflicting data

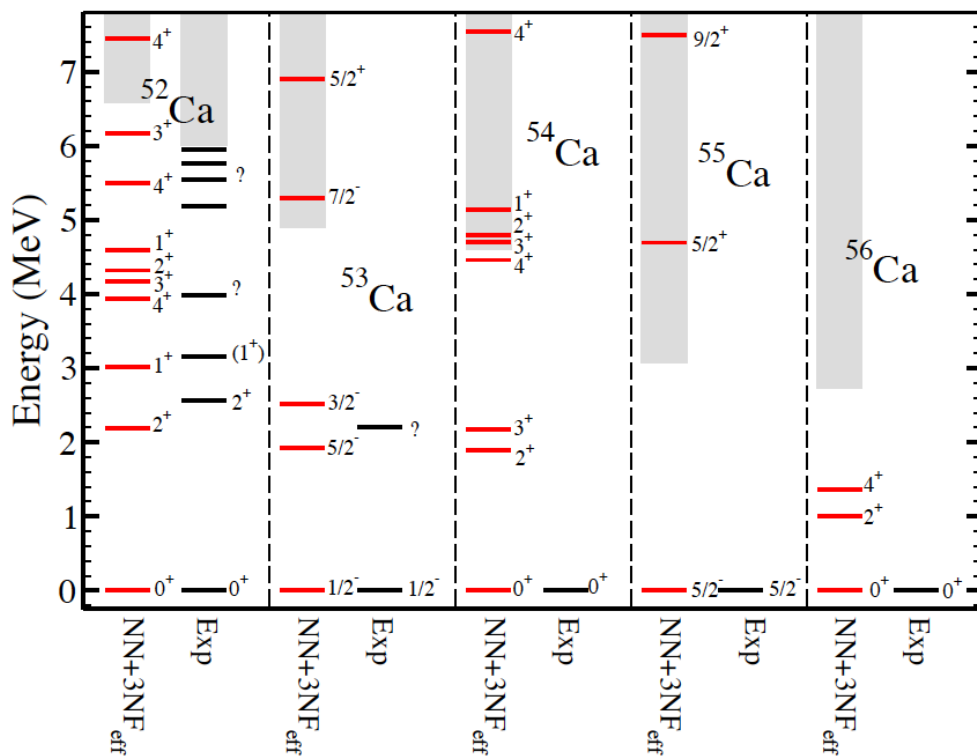
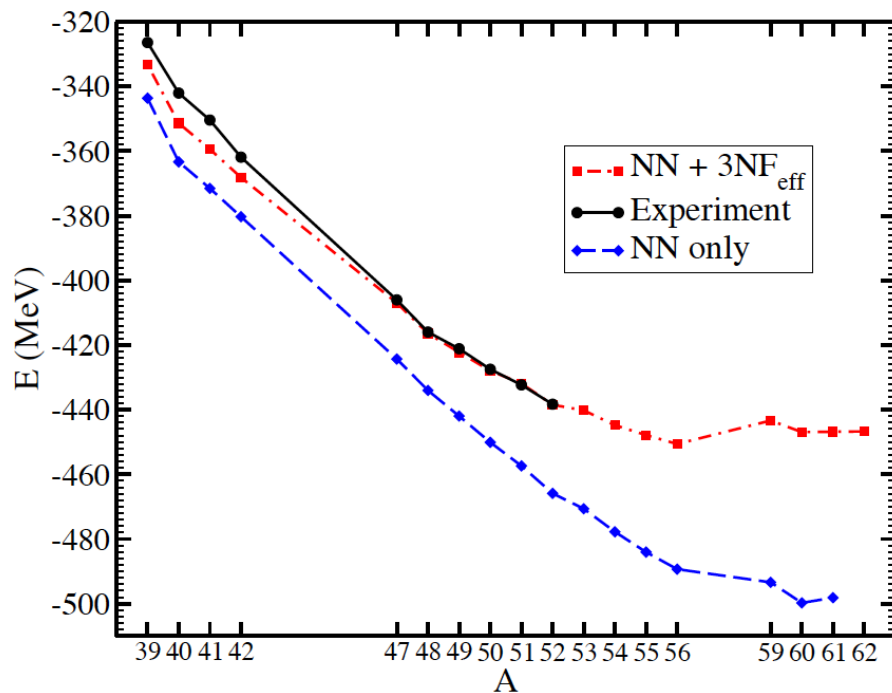


- The n - ${}^3\text{H}$ elastic cross section for 14 MeV neutrons, important for NIF, was not known precisely enough.
- Delivered evaluated data with required 5% uncertainty and successfully compared to measurements using an Inertial Confinement Facility
- “First measurements of the differential cross sections for the elastic n - ${}^2\text{H}$ and n - ${}^3\text{H}$ scattering at 14.1 MeV using an Inertial Confinement Facility”, by J.A. Frenje *et al.*, Phys. Rev. Lett. **107**, 122502 (2011)
<http://physics.aps.org/synopsis-for/10.1103/PhysRevLett.107.122502>

Larger Nuclei: Coupled-cluster method

Description of medium-mass open nuclear systems

G. Hagen et al., *Phys. Rev. Lett.* 109, 032502 (2012)




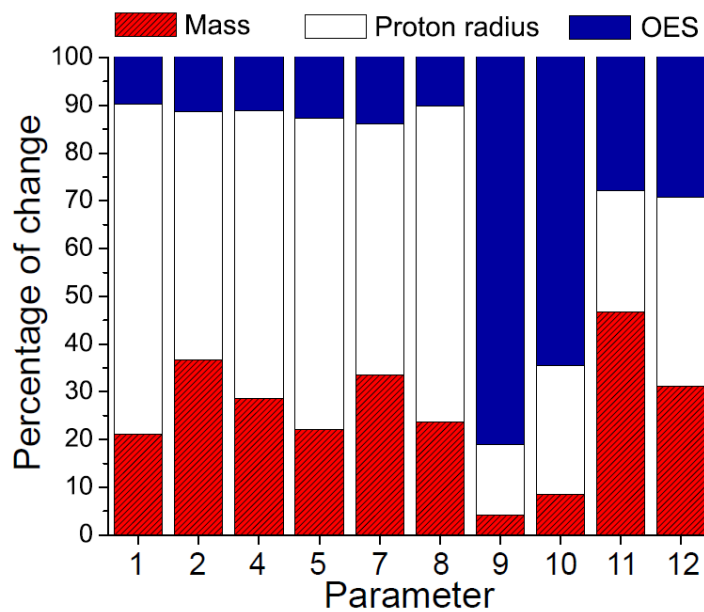

J^π	^{53}Ca		^{55}Ca		^{61}Ca	
	$\text{Re}[E]$	Γ	$\text{Re}[E]$	Γ	$\text{Re}[E]$	Γ
$5/2^+$	1.99	1.97	1.63	1.33	1.14	0.62
$9/2^+$	4.75	0.28	4.43	0.23	2.19	0.02

$1/2^+$ virtual state

- Strong coupling to continuum for neutron rich calcium isotopes
- Level ordering of states in the *gds* shell is contrary to naive shell model picture

Improved Density Functionals Neutron Drops, Masses, Fission,...

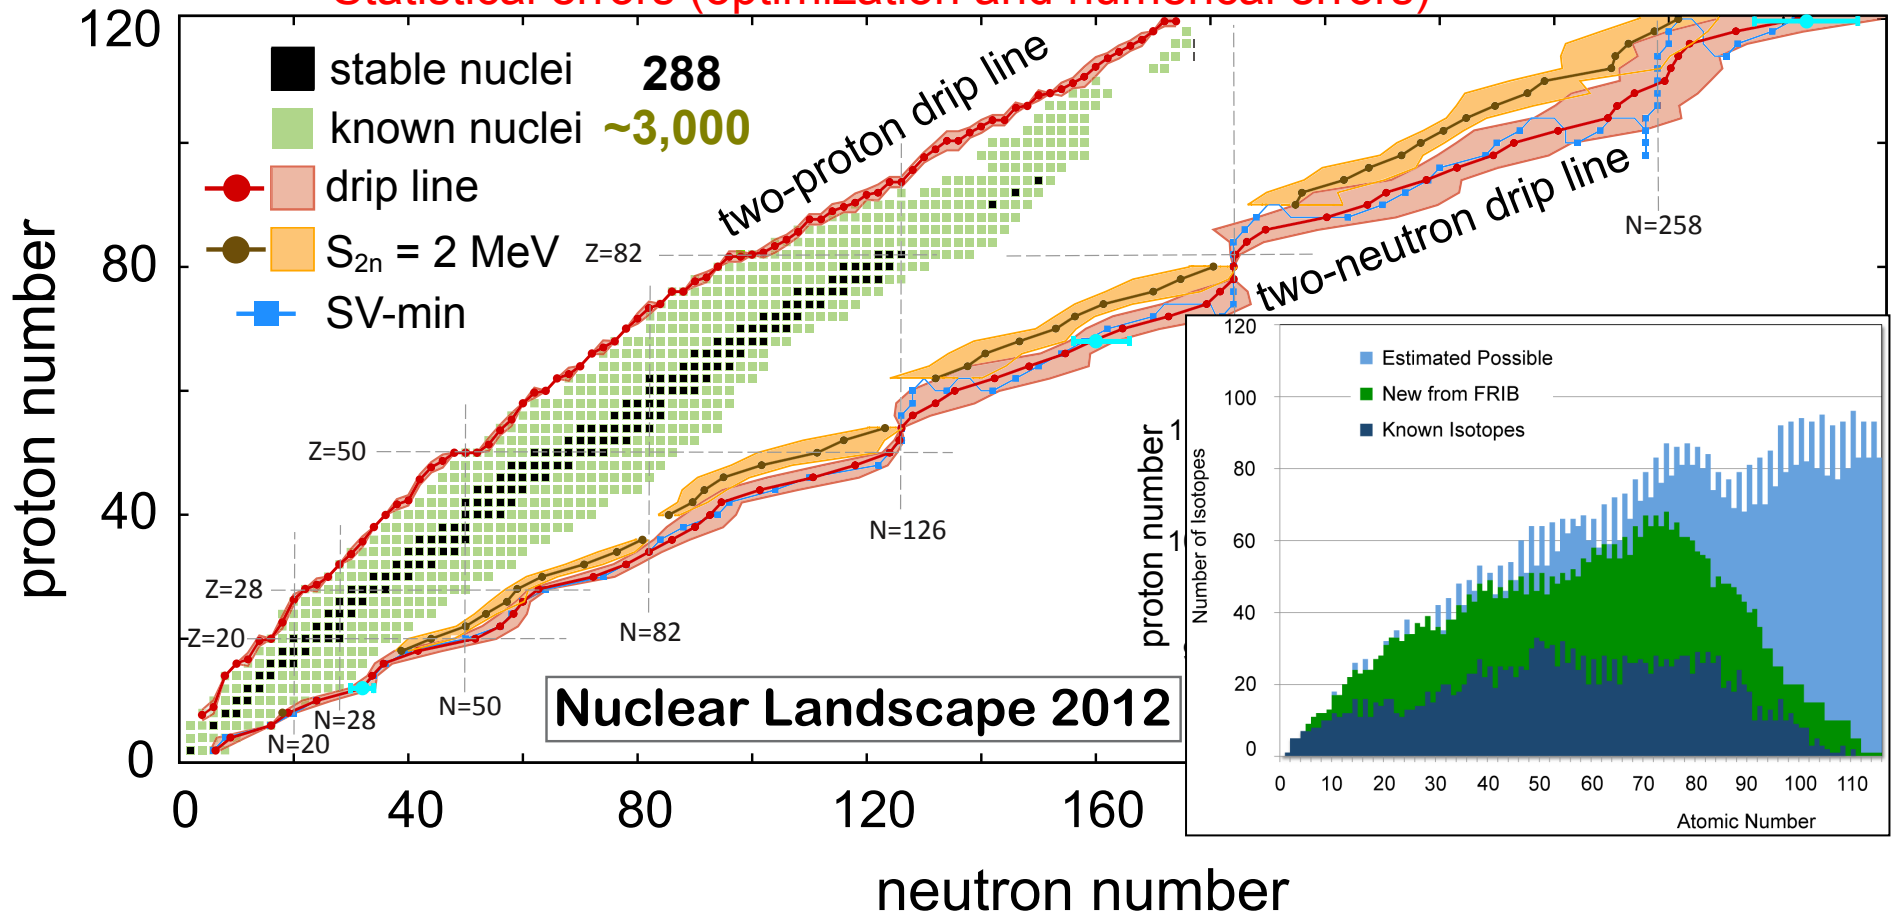
Derivative-free optimization, uncertainty quantification



<http://www.deixismagazine.org/2011/03/cracking-up-the-speed-of-dft/>
<http://www.deixismagazine.org/2011/03/pounding-out-atomic-nuclei/>
<http://www.mcs.anl.gov/news/detail.php?id=720>

The limits: Skyrme-DFT Benchmark 2012

- Systematic errors (due to incorrect assumptions/poor modeling)
- Statistical errors (optimization and numerical errors)



How many protons and neutrons can be bound in a nucleus?

Skyrme-DFT: $6,900 \pm 500_{\text{syst}}$

Erler et al.,
Nature 486, 509 (2012)

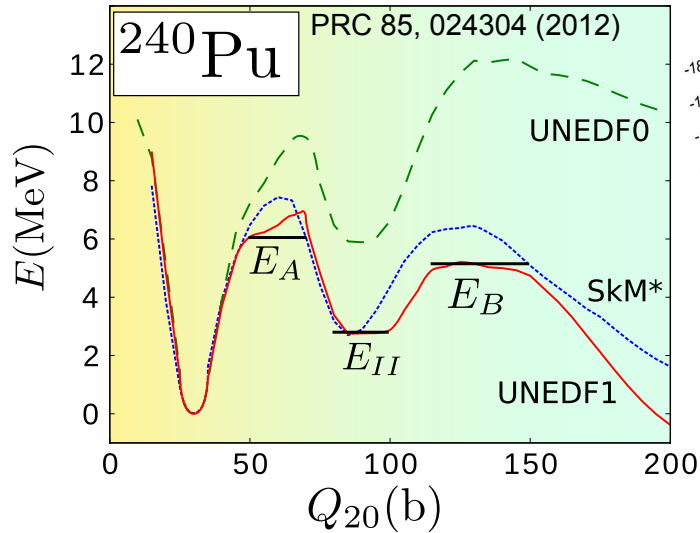
<http://www.livescience.com/21214-atomic-nuclei-variations-estimate.html>
<http://www.sciencedaily.com/releases/2012/06/120627142518.htm>

LACM, Fission: the ultimate challenge

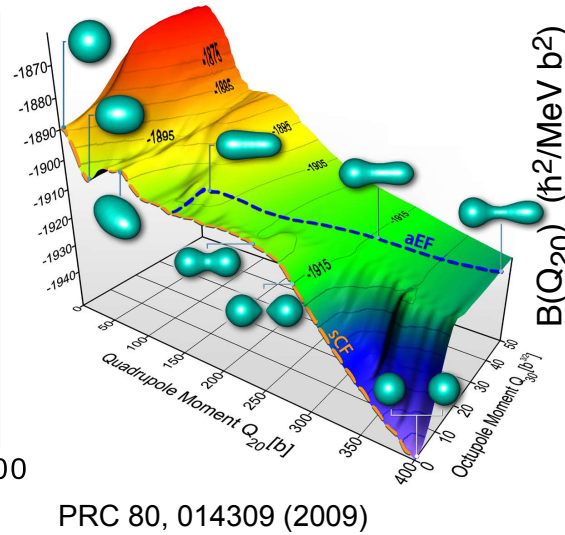
Stability of the heaviest nuclei, r-process, advanced fuel cycle, stockpile stewardship...



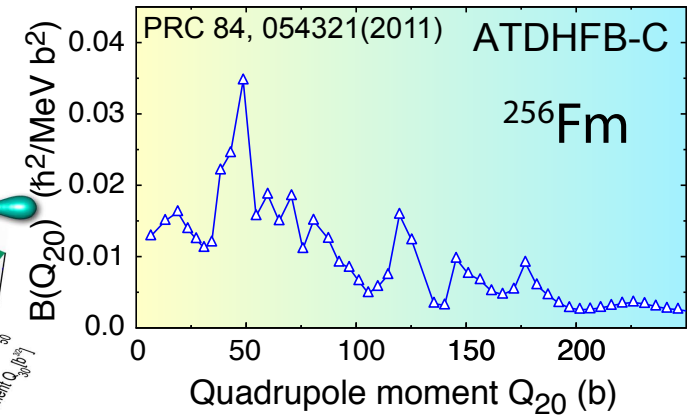
Optimized Functionals



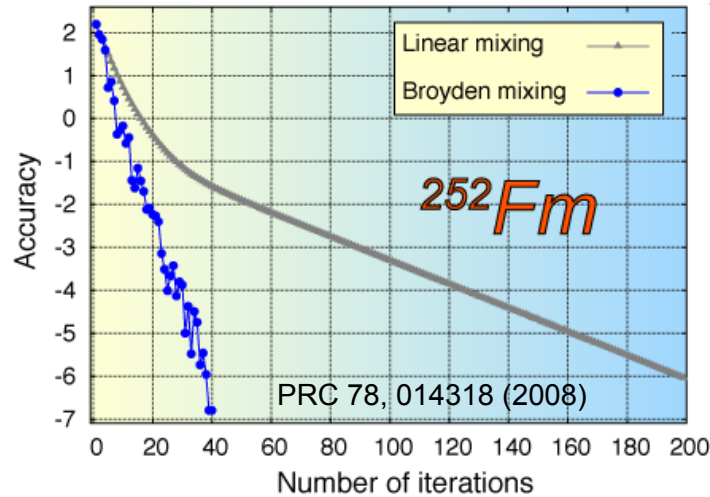
Large-scale DFT



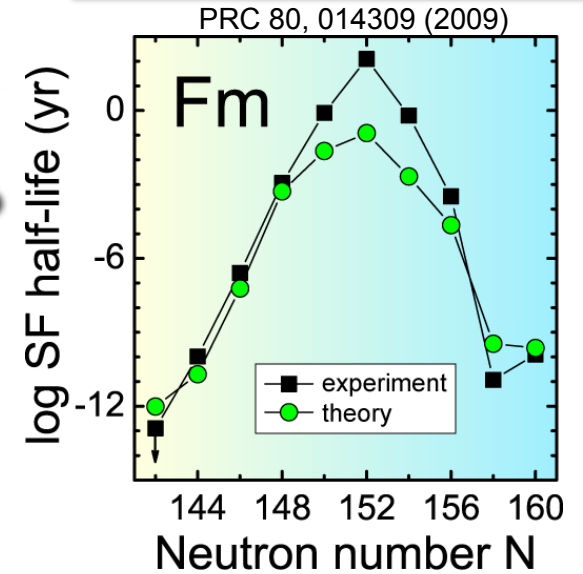

Collective dynamics




Numerical Techniques




Confrontation with experiment; predictions





Two DOE-sponsored workshops on exascale computing for nuclear physics, one common conclusion:


The microscopic descriptions of nuclear fusion and fission are our Priority Research Directions.



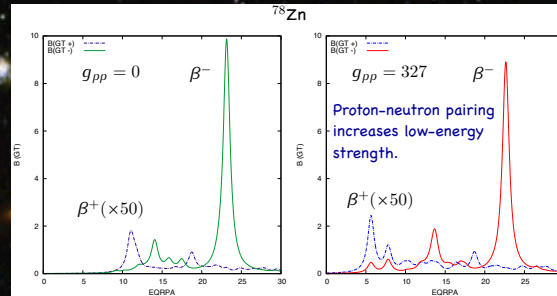
“The ultimate outcome of the nuclear fission project is a treatment of many-body dynamics that will have wide impacts in nuclear physics and beyond. The computational framework developed in the context of fission will be applied to the variety of phenomena associated with the large amplitude collective motion in nuclei and nuclear matter, molecules, nanostructures and solids”

“Fusion reactions in light nuclei are critical for both basic science and NNSA applications. These reactions aid in understanding the early universe and provide the energy that powers the stars. These reactions are also essential elements of NNSA science..”




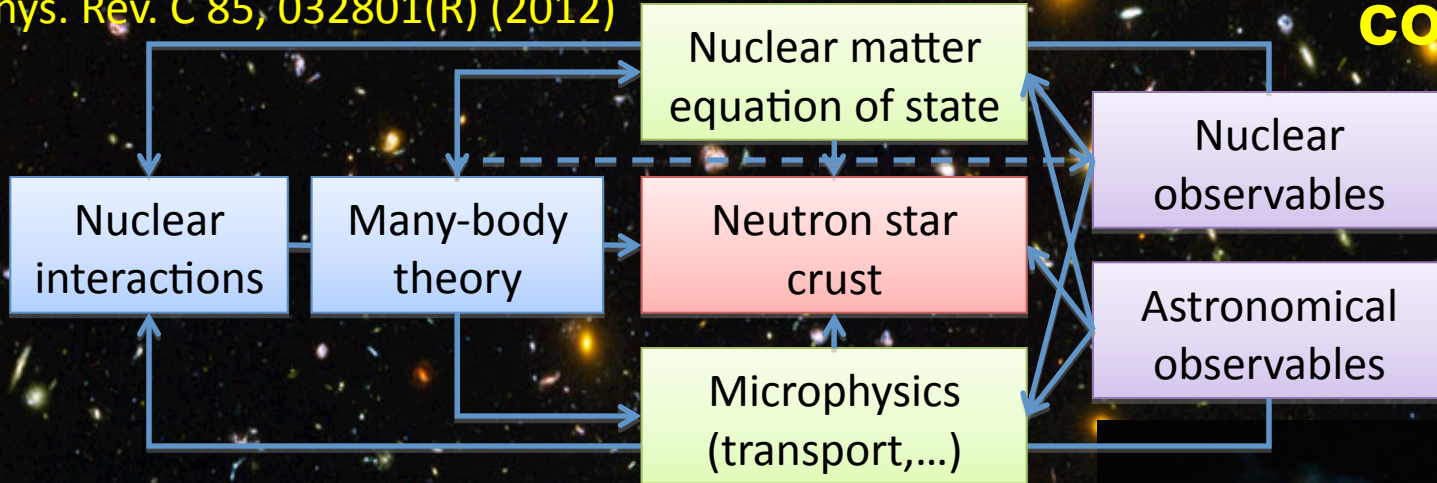
Quest for understanding the neutron-rich matter on Earth and in the Cosmos





Phys. Rev. C 85, 032801(R) (2012)



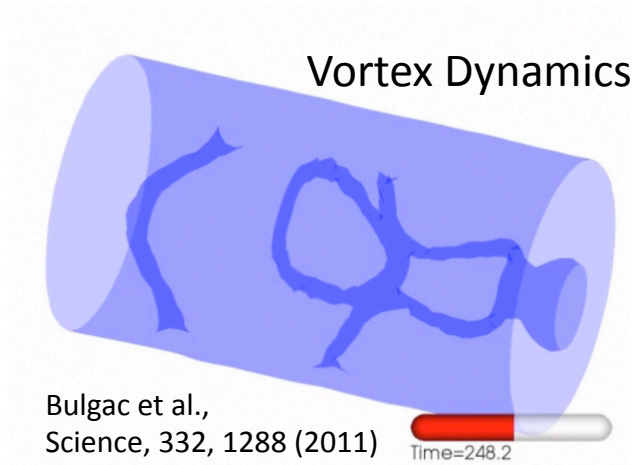
**RNB
Facilities/
NuN topical
collaboration**




Phys. Rev. Lett. 108, 081102 (2012)



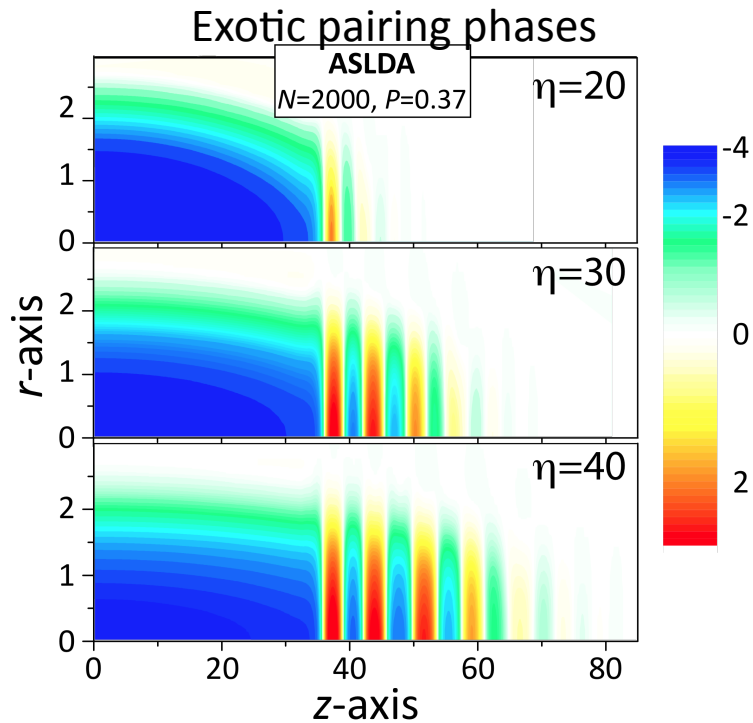
Connections to Other Fields: Cold Atoms




http://www.physicstoday.org/resource/1/phtoad/v64/i8/p19_s1



Equation of State



J. Pei et al., Phys. Rev. A 82, 021603(R) (2010)



Gezerlis and Carlson, Phys. Rev. C 77, 032801(R) (2008)
Carlson, Gandolfi, Gezerlis, PTEP (2012)

How many neutrons, protons can get along? Maybe 7,000

New study comes closer than ever to finding answer by estimating number of variations that can exist in an atom

Jump to discuss
Loading comments...


Recommend 29
Tweet 1
Share 13

By Clara Moskowitz

updated 6/27/2012 3:47:54 PM ET
Print | Font: A A + -

Scientists have long wondered whether there is a limit to the number of protons and neutrons that can be clustered together to form the nucleus of an atom. A new study comes closer than ever to finding the answer by estimating the total number of nucleus variations that can exist.

The periodic table of elements includes 118 known species of atoms, and each of these exists (either naturally or synthetically) in several versions with differing numbers of neutrons, giving rise to a total of about 3,000 different atomic nuclei. As technology has



Andy Spies / Oak Ridge National Laboratory

Pounding out atomic nuclei

Mike May | March 7th, 2011 | Updated: March 16th, 2011
Email This Post Print Share This


FILED UNDER Argonne,

Thousands of tiny systems called atomic nuclei – specific combinations of protons and neutrons – prove extremely difficult to study but have big implications for nuclear stockpile stewardship. To describe all of the nuclei and the reactions between them, a nationwide collaboration is devising powerful algorithms that run on high-performance computers.

Nuclear reactions, from fission in reactors to fusion in stars, depend on interactions between protons and neutrons that are building blocks of atomic nuclei.

Describing all of the nuclei and the reactions between them, however, demands powerful algorithms running on high-performance computers.

The Universal Nuclear Energy Density Functional (UNEDF) collaboration, which was created by the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program, focuses on developing such descriptions.



The UNEDF collaboration includes researchers from seven national laboratories – Ames, Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, and Pacific Northwest – and nine universities: Central Michigan, Iowa State, Michigan State, Ohio State, San Diego State, North Carolina at Chapel Hill, Tennessee-Knoxville, Texas A&M in Commerce and University of Washington. Recently, researchers in this collaboration made a significant advance through the use of density functional theory (DFT).

You are here: SC Home » News » Science Highlights Series » 2011 » Universities and DOE National Laboratories Join Forces to Understand the Nucleus of an Atom

News

Science Highlights Series

Press Releases

In the News

Science Highlights Series

2011

2010

2009

In Focus

Presentations & Testimony

Recovery Act

CONTACT US

Office of Science
U.S. Department of Energy
1025 Independence Ave., SW
Washington, DC 20585
P: (202) 586-5430
More Information >

03.28.11

Universities and DOE National Laboratories Join Forces to Understand the Nucleus of an Atom

Find out about this collaboration's efforts and accomplishments toward completing a portrait of the nuclear landscape.

Print SC RSS Feeds Text Size: A A A Share / Bookmark

Nuclear physics is the study of the tiny, massive core of an atom, a complex micro-world of particles and forces. Nearly all the mass in the visible universe is locked away in atomic nuclei, as is nearly all the energy. The physics of the nucleus lies at the heart of element formation in exploding stars, as well as sources of energy for public use and national defense.

Scientists strive for a comprehensive, unified description of all nuclei, a portrait of the nuclear landscape which incorporates all nuclear properties and forces in one framework. Such a model would allow for more accurate predictions of the nuclear reactions involved in all sorts of processes, from the creation of new elements to the improvement of nuclear reactors.

Around 50 researchers – theoretical physicists, computer scientists, and applied mathematicians – from nine U.S. universities, seven national laboratories, and research institutes across Europe and Japan, have come together in an effort to develop a more complete description of the atomic nucleus and its interactions. Their computational nuclear physics project, known as Universal Energy Density Functional (UNEDF), is led by Ewingusk (Argonne National Laboratory) and Witold Nazarewicz (University of Tennessee/Oak Ridge National Laboratory). The project is part of the U.S. Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) program, funded by the Office of Science.

A nucleus is one of the most complicated environments in nature because all fundamental forces come into play. The four fundamental forces are called the strong force, electromagnetic force, weak force, and gravity. The constituents of a nucleus are protons and neutrons (collectively referred to as nucleons), which are themselves made of fundamental particles known as quarks and gluons. Each




« Fluorescing diamonds inside living cells | Physics Update

Stirring superfluids

By Physics Today on June 13, 2011 10:52 AM | No TrackBacks

If you chill fermions enough, they can pair up to form bosons and settle into a single ground state, a Bose–Einstein condensate. In the case of helium-3 atoms, the resulting superfluid that flows without dissipation—provided the flow is not so energetic that it is apart or destroys the ground state's coherence. Until now, theorists could characterize in fermionic superfluids, but not the vigorous turbulence that results from shaking or stirring. In their simulations, Bulgac and his colleagues agitated a fermionic superfluid by shooting spherical projectiles through it or by stirring it with a laser beam. Turbulent superfluids are known to harbor tubes of quantized vorticity. As the figure below shows, the simulation how two vortex tubes (marked a and b) joined to form a ring, which then opens in a manner reminiscent of the unzipping of a DNA molecule during transcription. Bulgac's model of astronomers understand another agitated superfluid: the interior of a rapidly spinning (A. Bulgac et al., *Science* **332**, 1288, 2011.)—Charles Day



Science News

Physicists Pin Down Proton-Halo State

ScienceDaily (May 27, 2010) — A halo may be difficult to acquire in terms of virtue, but it can also be tough to calculate in terms of physics. Thomas Papenbrock, associate professor of physics and astronomy at the University of Tennessee, Knoxville, and his colleagues Gaute Hagen from Oak Ridge National Laboratory and Morten Hjorth-Jensen from the University of Oslo have managed to do just that, however, and report their findings in *Physical Review Letters*.

Home | Physics & Math | News

Quantum quirk makes carbon dating possible

15 July 2011 by David Shiga
Magazine issue 2821. [Subscribe and save](#)

RADIOCARBON dating relies on carbon-14 to decode an object's age, but the isotope has steadfastly refused to divulge the key to its own unusual longevity. The answer, it seems, lies in the bizarre rules of quantum physics.

Carbon-14 decays with a half-life of 5730 years, so it is often used to date objects up to about 50,000 years old (anything older would have negligible amounts of the stuff).

But most other atoms that decay in the same way - by converting one of their neutrons into a proton - disappear in less than a day. So what's different about carbon-14?

The nucleus of the carbon-14 isotope has six protons and eight neutrons. When it decays, one of the neutrons turns into a proton, and also releases an electron and a neutrino. The result is a nitrogen-14 nucleus with seven protons and seven neutrons.

International Impact of UNEDF

- JUSTIPEN (Japan)
 - FUSTIPEN (France)
 - FIDIPRO (Finland)
 - individuals worldwide
-
- Annual collaboration meetings of UNEDF with JUSTIPEN and FIDIPRO
 - Helping our Japanese colleagues to sharpen the case for nuclear theory supercomputing (K-machine)
 - Significant impact on low-energy nuclear theory effort worldwide (example: INPC 2010)
 - Unique worldwide (no other program has such a broad scope and such close PHY/CS/AM partnerships)

What makes this possible?

- Computational Resources
- Physicists (particularly early career are critical)
- Math/CS research and people (Lusk)

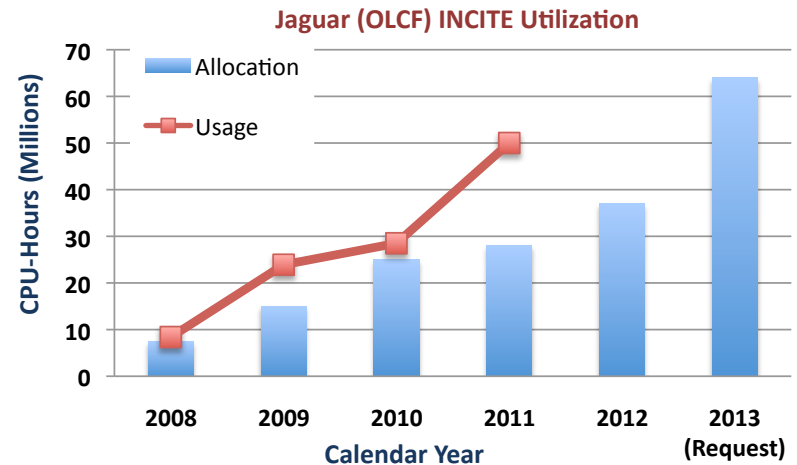


UNEDF Leadership-class computing

◆ SciDAC collaboration between applied mathematicians, computer scientists, and nuclear physicists leads to efficient utilization of leadership-class computing resources for nuclear physics problems

◆ Significant accomplishments, achieved through leadership-class computing

- Ab-initio calculations of C-12
- Understanding long lifetime of C-14
- Microscopic calculations of select medium-mass nuclei
- Improved energy-density functional UNEDF1

◆ 60% to 80% of computing resources used at leadership-class scale



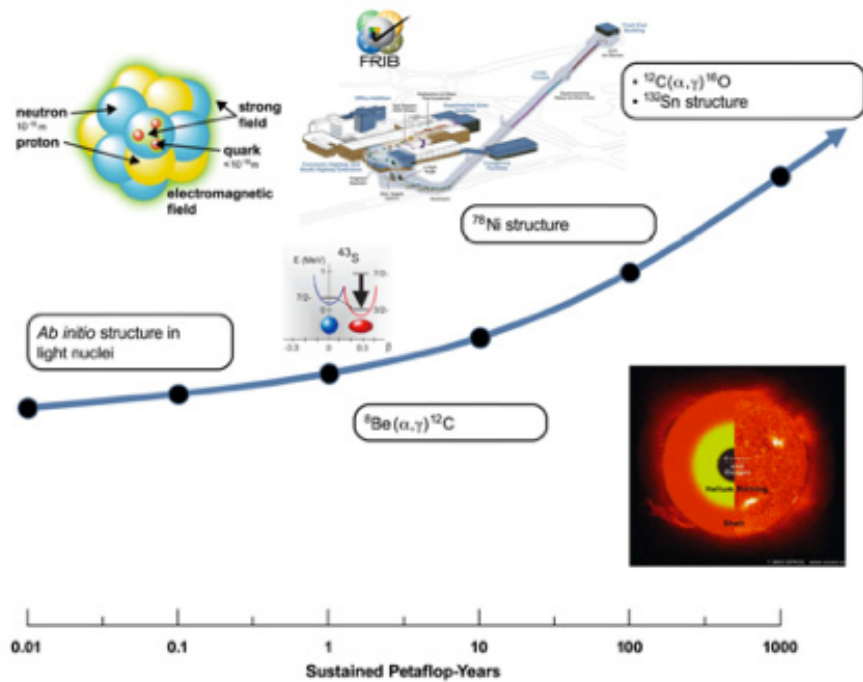
UNEDF SciDAC Collaboration
Universal Nuclear Energy Density Functional

Total Resource Usage / Request / Needs

	2009	2010	2011	2012 (alloc)	2013 (request)
INCITE					
Intrepid	12	9	14	19	45
Jaguar	22	28	50	35	65
NERSC	2	3	3	4	10
Jaguar	20	25	35	50	100
LLNL, LANL, ANL, ORNL	20	25	25	30	40
TOTAL	70	90	120	140	250

In million core-hours (approximate)

Nuclear Reactions



2012

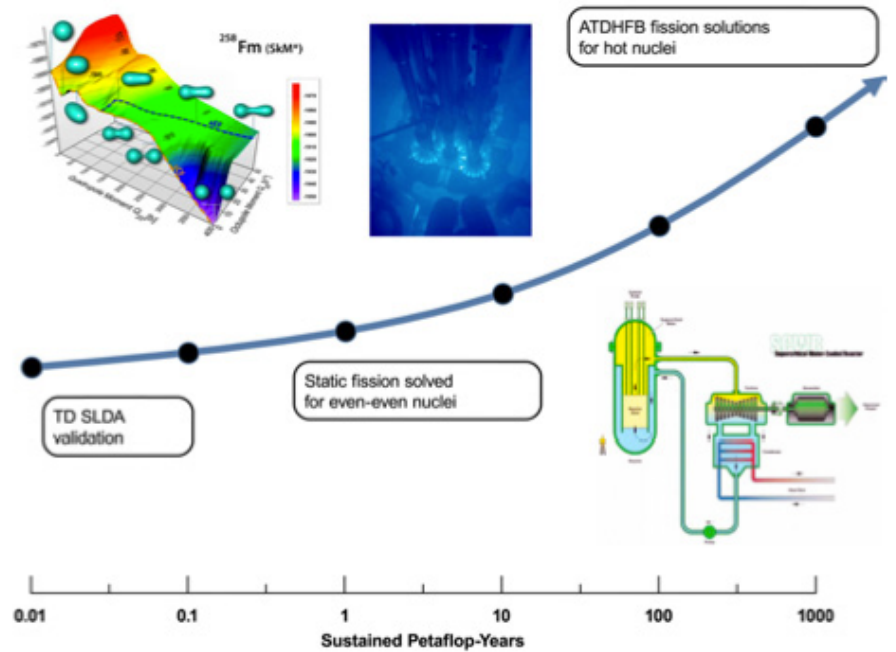


2020 2020



2030 2028 2020

Microscopic Theory of Fission



2012



2020 2020




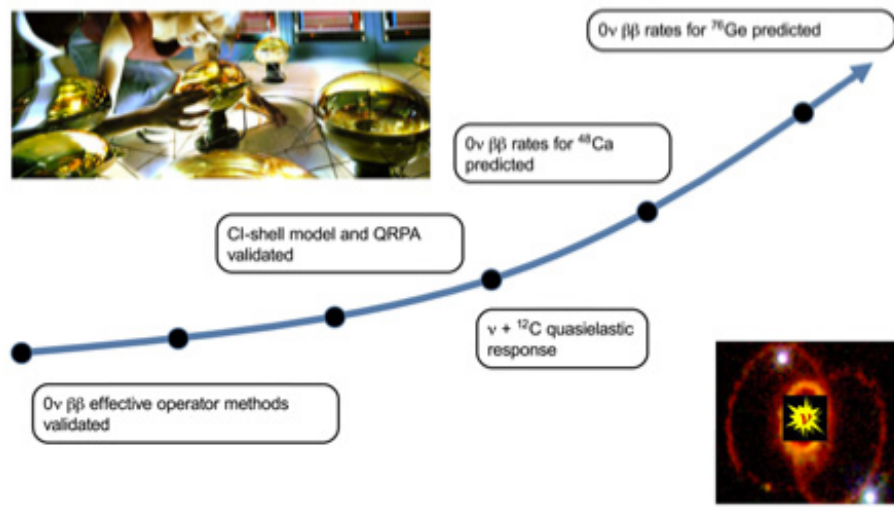
2030 2028 2020

Desired Trajectory


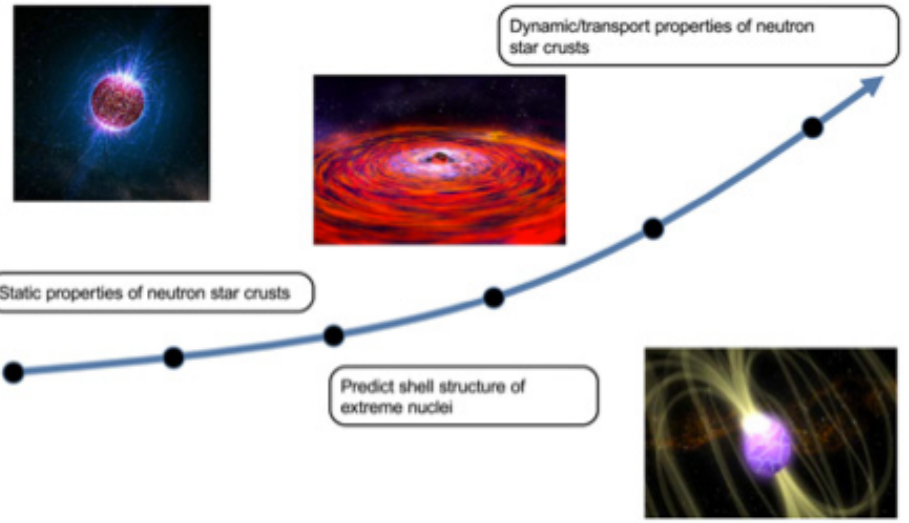
Flat Trajectory

Flat-Flat Trajectory

Neutrinos as Nuclear Physics Laboratories



Physics of Extreme Neutron-Rich Nuclei and Neutron Stars



Desired Trajectory
 Flat Trajectory
 Flat-Flat Trajectory

Junior Scientists in UNEDF

POST-DOCTORAL ASSOCIATES (2010)

Christopher Calderon, LBNL (staff, Numerica co.)
Joaquin Drut (Professor, UNC)
Stefano Gandolfi, LANL (staff, LANL)
Kai Hebel, OSU (TRIUMF)
Heiko Hergert, MSU (OSU)
Jason Holt, UTK/ORNL
Eric Jurgenson, LLNL (staff, LLNL)
Markus Kortelainen, UTK (U. Jyväskylä)
Plamen Krastev, UCSD (research, Harvard)
Pieter Maris, ISU (Research Prof. ISU)
Eric McDonald, MSU (staff scientist, MSU)
Gustavo Nobre, LLNL (BNL, NNDC)
Junchen Pei, UTK (Prof., Pekin U.)
Nicolas Schunck UTK (staff, LLNL)
Roman Senkov, CMU
Ionel Stetcu, UW (staff, LANL)
Jun Terasaki, UNC (staff, U. Tsukuba)
Stefan Wild, ANL (staff, ANL)

2010: Early
Career Award

2011: Faculty
UNC/Chapel
Hill

2012: Faculty
Guelph

Relevant instruction (workshops, courses)
is crucial for the future of the field

Effect of UNEDF on workforce

Year-1: 9 students, 17 postdocs;

Year-2: 12 and 12;

Year-3: 10 and 18;

Year-4: 11 and 19



2010:
Math/CS
Staff ANL



2010: Staff
LLNL



2012:
Harvard
Research
Computing



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Mathematics and Computer Science in Two Nuclear Physics SciDAC Projects

UNEDF: 2007-2011

NUCLEI: 2012-2016 (hopefully)



SciDAC

Scientific Discovery through Advanced Computing



Outline

- The UNEDF project just ending
 - goals
 - structure
 - Math/CS areas
- Some Specific accomplishments
- The NUCLEI project just starting (no logo yet)
- Lessons learned



Universal Low-Energy Density Functional (UNEDF)

- Sponsors: NP, ASCR, and NNSA
- Math/CS goals of this SciDAC-2 project:
 - Improve algorithms and software in support of moving codes (existing, for the most part) to petascale computing facilities, thus enabling new scientific results
 - Achieve results in mathematics and computer science applicable outside the project
- Structure:
 - Not a Math/CS team and a physics team
 - Rather, a collection of specific projects requiring math/CS and physics both, tied together by an overarching physics objective

Math/CS Areas

- Optimization
 - Load Balancing
 - Eigenvalues
 - Frameworks
 - Programming Models
-
- Each one part of a specific physics subproject of UNEDF

Computer Science in UNEDF

- Several partnerships; model collaborations
- Collaborations established before proposal was submitted

Green's Function Monte Carlo

Physics challenge: ab-initio description of 12C, 14C, 16O

Computational challenges: 800,000 core hours each on Argonne's IBM BlueGene/P for 12C. 14C is 60 times harder, and 16O is more than 1,000 times harder.

Exploiting BG/Q will require a redesign of the ADLB Load Balancing Library.

Coupled-cluster Code Suite

Physics challenge: ab-initio description of structure and reaction properties of medium-mass and heavy nuclei.

Computational challenges: iteratively solve coupled, non-linear equations scaling roughly as $O(N^4)$. Current calculations include $N \sim 1,000$ (oxygen and calcium isotopes) with calculations planned to $N \sim 7,000$ (nickel and tin isotopes). To address scaling issues and facilitate feature additions, a generalized tensor contraction engine (TCE) and global arrays will be implemented. Improved I/O, checkpoint and fault tolerance tools will be required, as well as adaptive mesh technologies for collision modeling.

No-Core Shell Model

Physics challenge: ab-initio description of light nuclei

Computational challenges: typically solve for the lowest 10-20 eigenvalues and eigenvectors of a large sparse symmetric matrix with dimensions up to 1010. A single case can run for about 5 hours of wallclock time using the full Jaguar PF machine ($\sim 240,000$ cores) including post processing. Future applications will involve significantly increased matrix dimensions.

Nuclear Density Functional Theory

Physics challenge: global description of nuclear properties and decays (including fission) throughout the nuclear landscape with the focus on medium-mass and heavy nuclei.

Computational challenges: huge systems (over 10^7) of integro-differential equations per configuration solved iteratively. This is usually done by projection into a pre-determined basis with subsequent computations employing mostly dense matrices. Various boundary conditions assumed leading to discrete, resonance, and scattering problems. A realistic description of fission requires up to 10^9 configurations. Requirements: scalable IO; fault tolerance; on-the-fly simulation steering; high-dimensional array manipulations via tensor contractions; novel use of visualization tools for analysis. The use of adaptive multi-resolution techniques such as MADNESS framework will be required to solve extreme problems. MADNESS parallel runtime ensures efficient execution on leadership computers.

Nuclear Density Functional Optimizations

Physics challenge: Simulation-based, derivative-free optimization involving 10-20 parameters; sensitivity analysis for uncertainty and correlation estimations.

Computational challenges: Efficient parameter estimation and multilevel optimization algorithms for large scale (many core) DFT calculations. Estimating derivatives of noisy simulations.


Time-dependent Density Functional Theory

Physics challenge: Dissipative nuclear processes - nuclear fission and fusion, nuclear reactions, and time-dependent phenomena in nuclei and related systems (cold gases, neutron stars)


Computational challenges: Determination of nuclear ground states requires full diagonalization of large matrices (10^6) and higher) - one diagonalization requires the whole JaguarPF for several hours today; efficient I/O for large data volumes for check point and restart; in situ analysis for reduction of large data and novel use of visualization tools for analysis. A real time stochastic formulation of the current approach will require new communication algorithms for reductions and overlaps, fault tolerance, in situ data reductions, likely higher precision numerics, and improved I/O.

Derivative-free Optimization for Energy Density Functional Calibration

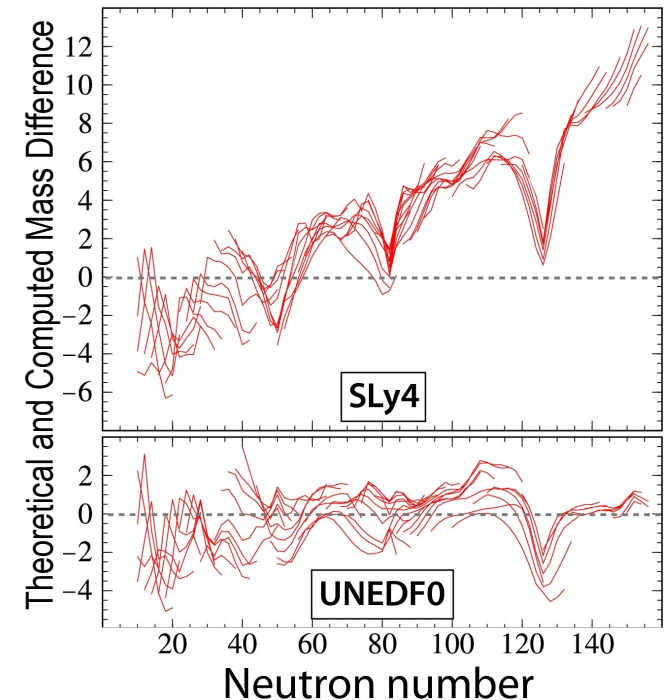
New algorithm for tuning large-scale nuclear structure simulations turns days into hours



- Energy density functional (EDF) predictions rely on large-scale computer simulations that must be calibrated to experimental data
- TAO 2.0's **POUNDERS** developed for UNEDF to exploit the mathematical structure of this calibration problem
- Substantial computational savings over alternative algorithms enables fitting of complex EDFs



- Previous optimizations required too many evaluations to obtain desirable features exhibited by *UNEDF0*, *UNEDF1*, ...
- Derivative-free sensitivity analysis procedure developed for UNEDF exposes correlations and constraining data in 1 minute using 20k cores
 - “Nuclear Energy Density Optimization” Kortelainen et al., *PhysRevC* ‘10.
 - “N.E.D.O.: Large Deformations” Kortelainen et al., *PhysRevC* ‘12.
 - “Occupation Number-based Energy Functional for Nuclear Masses” Bertolli et al., *PhysRevC* ‘12.




Optimal Derivatives of Noisy Numerical Simulations

Computational Noise

In all computations of DOE interest containing


- adaptivity,
- discretizations
- iterative methods
- petaflops,
- roundoff errors

❖ Includes *deterministic* computations



Noise Impacts in UNEDF & Beyond

- Uncertainty in computed outputs
- Unstable derivative estimates for sensitivity analysis
- Can be unrelated to/overwhelm truncation error
- Blurs relationship between tolerance values and stability



Tools & Techniques

- *ECNoise* provides reliable estimates of stochastic and deterministic noise in few simulations
- Nonintrusive **stability bounds for extreme scale simulations**, can instruct precision levels/tolerances for subroutines
- Optimal difference parameters calculated without computationally expensive parameter sweeps

Optimal step was obtained with only two simulation evaluations. Classical approach (circles): result of a sweep across 100 difference parameters, each point requiring a new simulation.

“Estimating Computational Noise,” Moré & Wild, *SIAM Sci. Comp.*, '11

“Estimating Derivatives of Noisy Simulations,” Moré & Wild, *ACM TOMS* '12

Making a simple programming model scalable: The Asynchronous Dynamic Load Balancing Library



Objectives

- Enable Green's Function Monte Carlo calculations for ^{12}C on full BG/P as part of UNEDF project
- Simplify programming model
- Scale to leadership class machines

Impact

- Demonstrate capabilities of simple programming models at petascale and beyond
- Show path forward with hybrid programming models in library implementation

Improved Efficiency (compute time/wall time) with more nodes



Progress

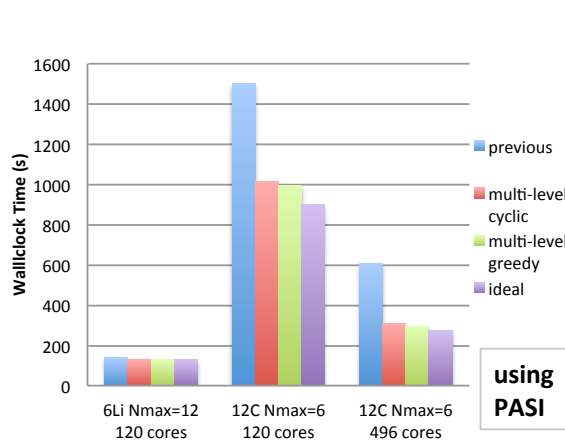
- Initial load balancing was of CPU cycles
- Next it became necessary to balance memory utilization as well
- Finally ADLB acquired the capability to balance message flow
- “More Scalability, Less Pain” by E. Lusk, S.C. Pieper and R. Butler published in SciDAC Review 17, 30 (2010)

MFDn: Total-J Progress

- ❖ **M-scheme approach:** works directly on the Hamiltonian, extracts all low energy states
- ❖ **J-scheme approach:** alternative to find a large number of low energy states for a prescribed total angular momentum (J) value
- ❖ **Targeted applications:** investigation of nuclear level densities, evaluation of scattering amplitudes
- ❖ **Total-J code:** implementation of the J-scheme approach in Fortran, MPI
- ❖ **CS Challenges in Total-J:** Three distinct phases, each with very different computing and storage characteristics


Phase 1: Construction of the J-basis

Implemented a multi-level greedy load balancing algorithm



Phase 2: Invariant Subspace Projection


In-core implementation to reduce I/O overheads in the out-of-core version



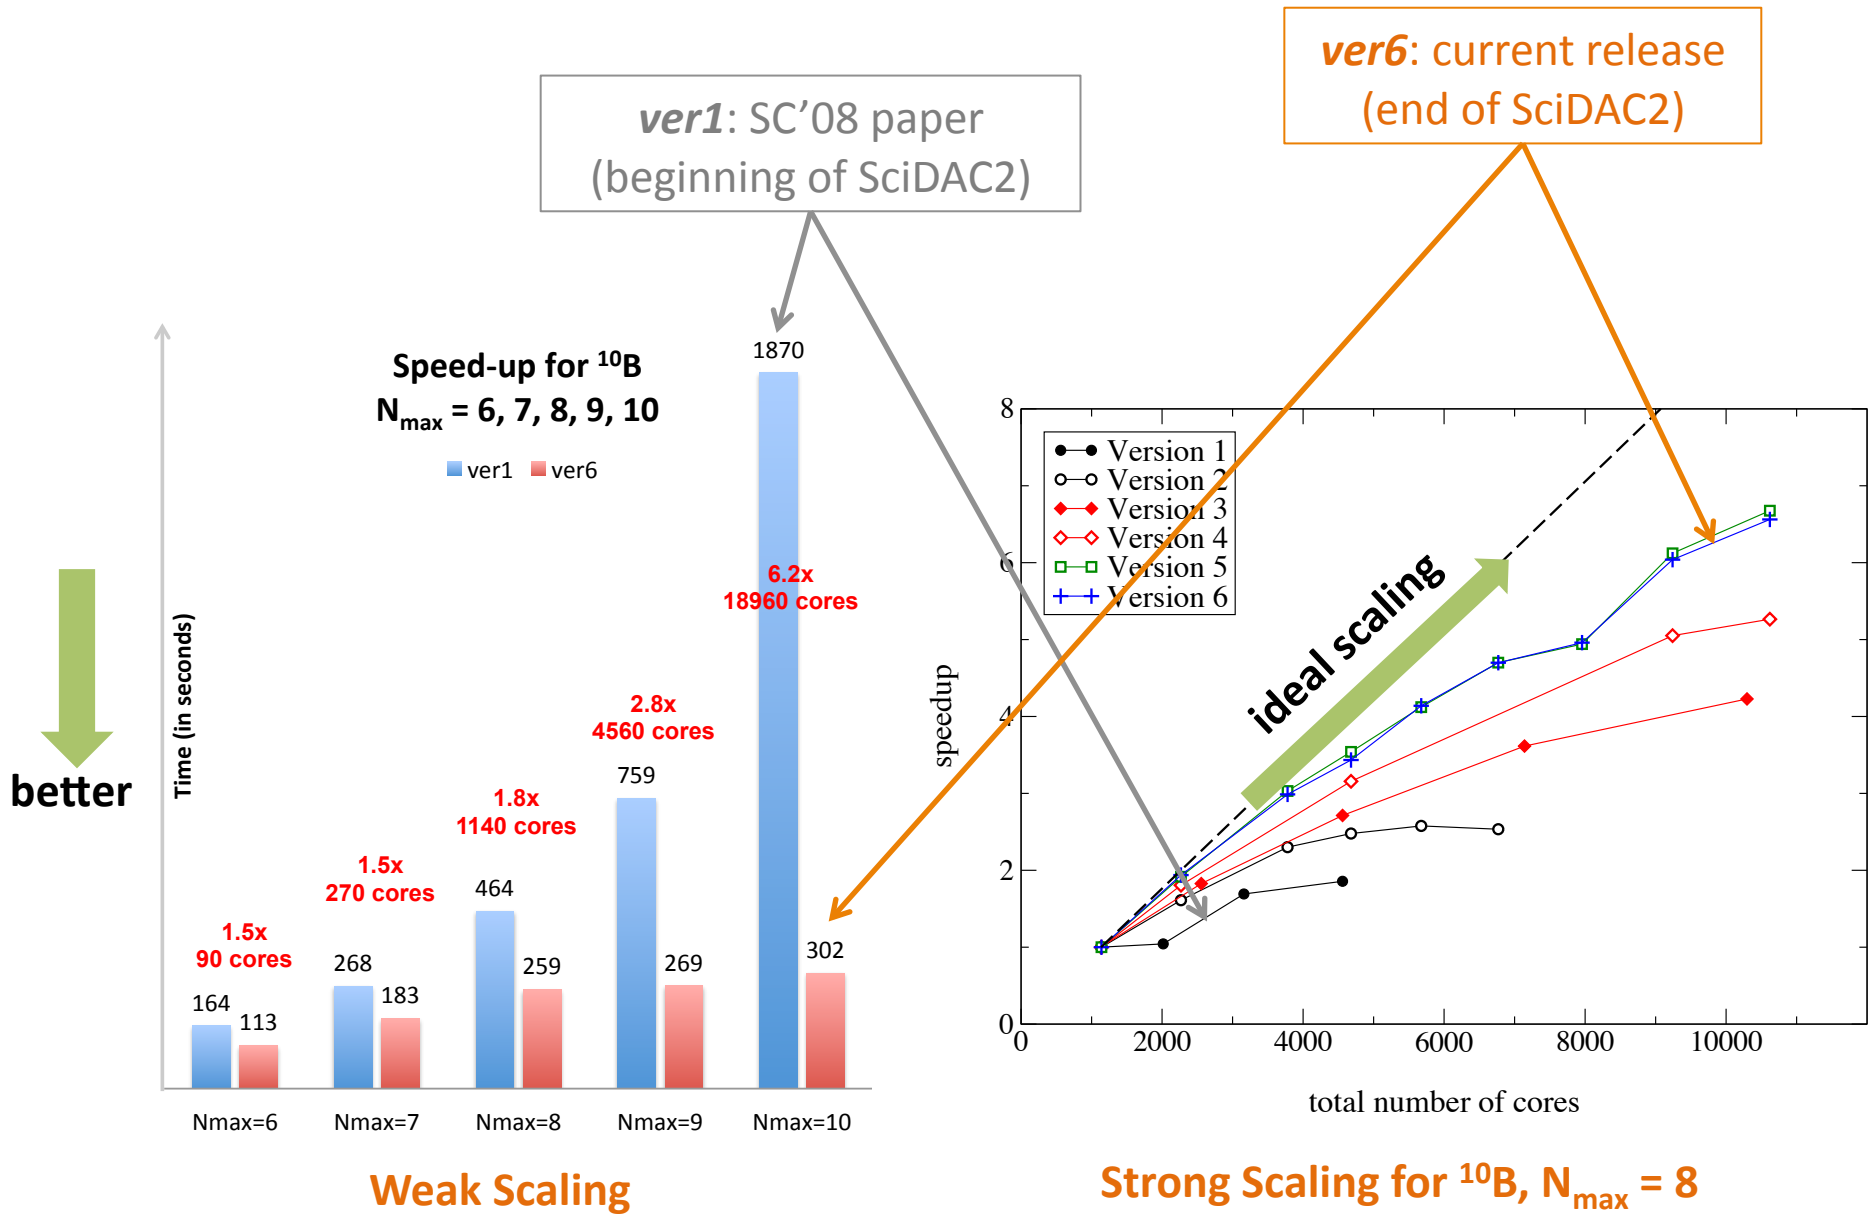
${}^6\text{Li}$, $N_{\max}=14$, $J=3$
 dimension of $H \sim 1.7 \times 10^8$
10x speed-up, default striping
4x speed-up, optimal striping

Application:

Predicting the Nuclear Level Density of ${}^6\text{Li}$ (PRELIMINARY)



M-scheme: Performance Improvements



One Demonstration of NP – ASCR Coupling: Over 20 joint publications resulting from UNEDF

- "Real-Time Dynamics of Quantized Vortices in a Unitary Fermi Superfluid," A. Bulgac, Y.-L. Luo, P. Magierski, K.J. Roche, and Y. Yu, *Science* 332, 1288 (2011).
- "Origin of the anomalous long lifetime of ^{14}C ," P. Maris, J. P. Vary, P. Navratil, W. E. Ormand, H. Nam, and D. J. Dean, *Phys. Rev. Lett.* 106, 202502 (2011).
- "More Scalability, Less Pain," E. Lusk, S.C. Pieper and R. Butler, *SciDAC Review* 17, 30 (2010).
- "Nuclear Energy Density Optimization," M. Kortelainen, T. Lesinski, J. More, W. Nazarewicz, J. Sarich, N. Schunck, M. V. Stoitsov, and S. Wild, *Phys. Rev. C* 82, 024313 (2010).
- "One-quasiparticle States in the Nuclear Energy Density Functional Theory," N. Schunck, J. Dobaczewski, J. McDonnell, J. More,, W. Nazarewicz, J. Sarich, and M.V. Stoitsov, *Phys. Rev. C* 81, 024316 (2010).
- "Scaling of ab-initio nuclear physics calculations on multicore computer architectures," P. Maris, M. Sosonkina, J. P. Vary, E. G. Ng and C. Yang, *International Conference on Computer Science, ICCS 2010, Procedia Computer Science* 1, 97 (2010).
- "Hamiltonian light-front field theory in a basis function approach," J. P. Vary, H. Honkanen, Jun Li, P. Maris, S. J. Brodsky, A. Harindranath, G. F. de Teramond, P. Sternberg, E. G. Ng, C. Yang, *Phys. Rev. C* 81, 035205 (2010).
- "Ab initio nuclear structure: The Large sparse matrix eigenvalue problem," J.P. Vary, P. Maris, E. Ng, C. Yang, and M. Sosonkina, *J. Phys. Conf. Ser.* 180, 012083 (2009).
- "Fast Multiresolution Methods for Density Functional Theory in Nuclear Physics," G. I. Fann, J. Pei, R. J. Harrison, J. Jia, J. Hill, M. Ou, W. Nazarewicz, W. A. Shelton, and N. Schunck, *Journal of Physics: Conference Series* 180, 012080 (2009).
- "Solution of the Skyrme-Hartree-Fock-Bogolyubov equations in the Cartesian deformed harmonic-oscillator basis. (VI) HFODD (v2.38j): a new version of the program," J. Dobaczewski, W. Satula, B.G. Carlsson, J. Engel, P. Olbratowski, P. Powalowski, M. Sadziak, J. Sarich, N. Schunck, A. Staszczak, M. Stoitsov, M. Zalewski, H. Zdunczuk, *Comp. Phys. Comm.* 180, 2361 (2009).
- "Towards The Universal Nuclear Energy Density Functional," M. Stoitsov, J. More, W. Nazarewicz, J. C. Pei, J. Sarich, N. Schunck, A. Staszczak, and S. Wild, *Journal of Physics: Conference Series* 180, 012082 (2009).
- "Deformed Coordinate-Space Hartree-Fock-Bogoliubov Approach to Weakly Bound Nuclei and Large Deformations," J.C. Pei, M.V. Stoitsov, G.I. Fann, W. Nazarewicz, N. Schunck and F.R. Xu, *Phys. Rev. C* 78, 064306 (2008).
- "Time-Dependent Density Functional Theory Applied to Superfluid Nuclei," A. Bulgac and K.J. Roche, *J. Phys.: Conf. Ser.* 125, 012064 (2008).
- ...

A SciDAC-3 Project: Nuclear Computational Low Energy Initiative (NUCLEI)

- Mostly same cast as UNEDF
- Targeting next generation architectures on path to exascale
- Examples of specific plans in Math/CS...

Optimization Plans in NUCLEI

New Optimization Capabilities

- State-of-the art mathematical/numerical optimizations of next-generation EDFs (*with LLNL, ORNL, Tennessee*)
- Optimization of basis states and non-perturbative coupling constants arising in chiral Hamiltonians (*with ISU*)
- Enable nucleus lifetime computations with collective action minimizations (*with LLNL, others*)
- Exploit additional parallelism at the simulation-optimization interface
- Extend POUNDERS to address missing states and available sensitivity information
- Incorporate uncertainties and QUEST technologies

Coupling NUCLEI Subgroups

- Incorporate new observables from various NUCLEI subgroups
 - giant resonance data,
 - binding energy of neutron droplets in a trap,
 - ...

Deploy code optimization tools

- Introduce performance, energy, and resilience tools developed by the SUPER SciDAC Institute for use in NUCLEI codes
- Deliver representative NUCLEI computational kernels to SUPER



Future Plans for the Asynchronous Dynamic Load Balancing Library

- The Exascale target for GFMC (Green's Function Monte Carlo) code is ^{16}O .
- The overall structure of GFMC can remain the same if
 - Effective use of many-core nodes can be made
 - may need to go beyond current OpenMP approach
 - ADLB can support multi-node parallelism on single work units
- This will require major changes to ADLB that take advantage of new features in MPI-3.

MFDn: Future Work

- Version 14 (under development)
 - Efficient construction of the Hamiltonian, faster SpMV
 - 3-body interaction support is underway
- MFDn on different architectures
 - Investigating the use of GPU (Titan) and BG/Q (Mira) platforms
 - Out-of-core implementation on SSD-equipped clusters
- Compression of the Hamiltonian to enable the study of heavier nuclei

Summary (for Math/CS)

- The structure of the UNEDF project was the key to its success
 - Math/CS//Physics partnerships formed locally at proposal-writing stage, and developed over the 5-year timeframe of the project
- Significant science resulted
- Future looks bright
 - Many hard steps now behind us
 - Collaborations quite deep at this point
 - Codes have been significantly improved for scalability
 - Many interesting steps ahead of us
 - Exascale provides new set of challenges

NUclear Computational Low-Energy Initiative

