Role of Supercomputers
in Basic Science Simulations

James P. Vary, Iowa State University

NTSE-2016, Pacific National University
Khabarovsk, Sept. 19-23, 2016

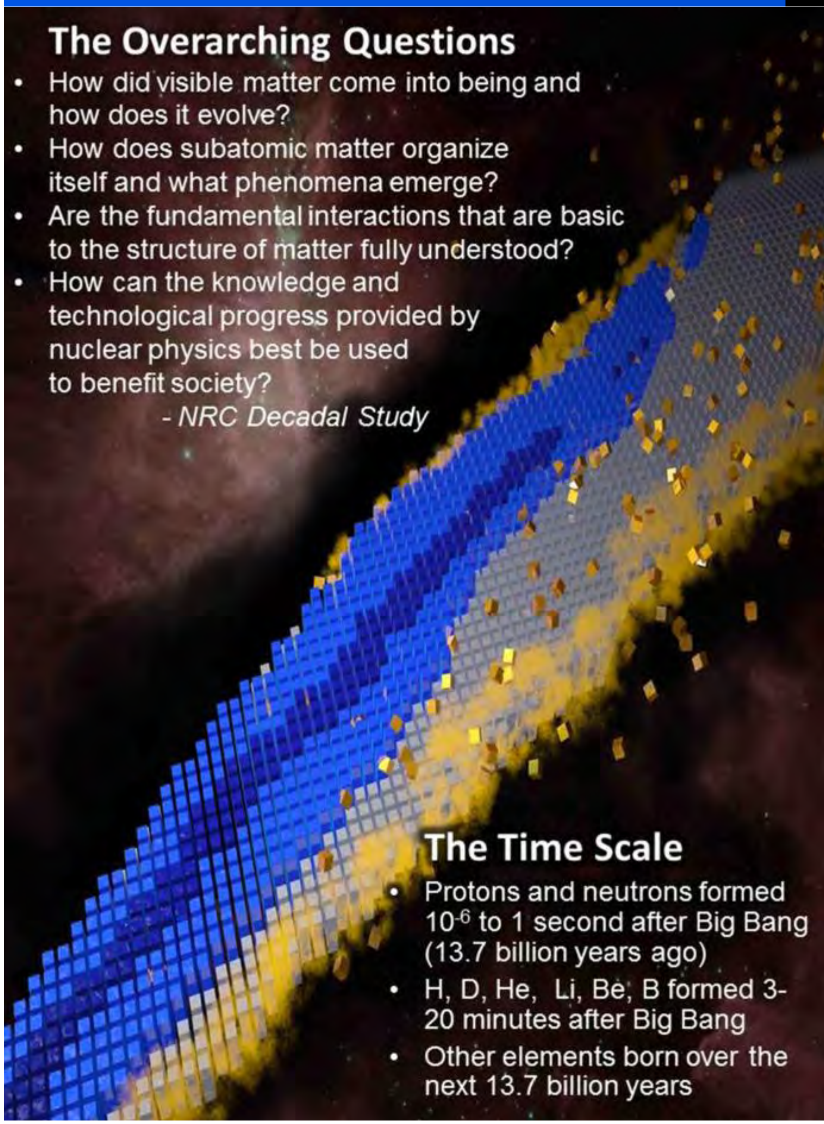

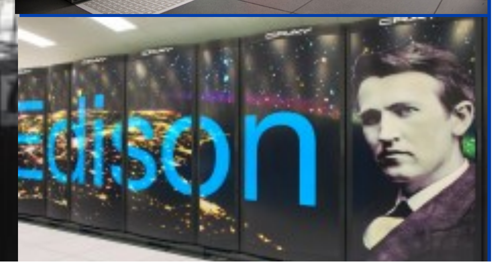

World's Top 500 Supercomputers

Projected Performance Development

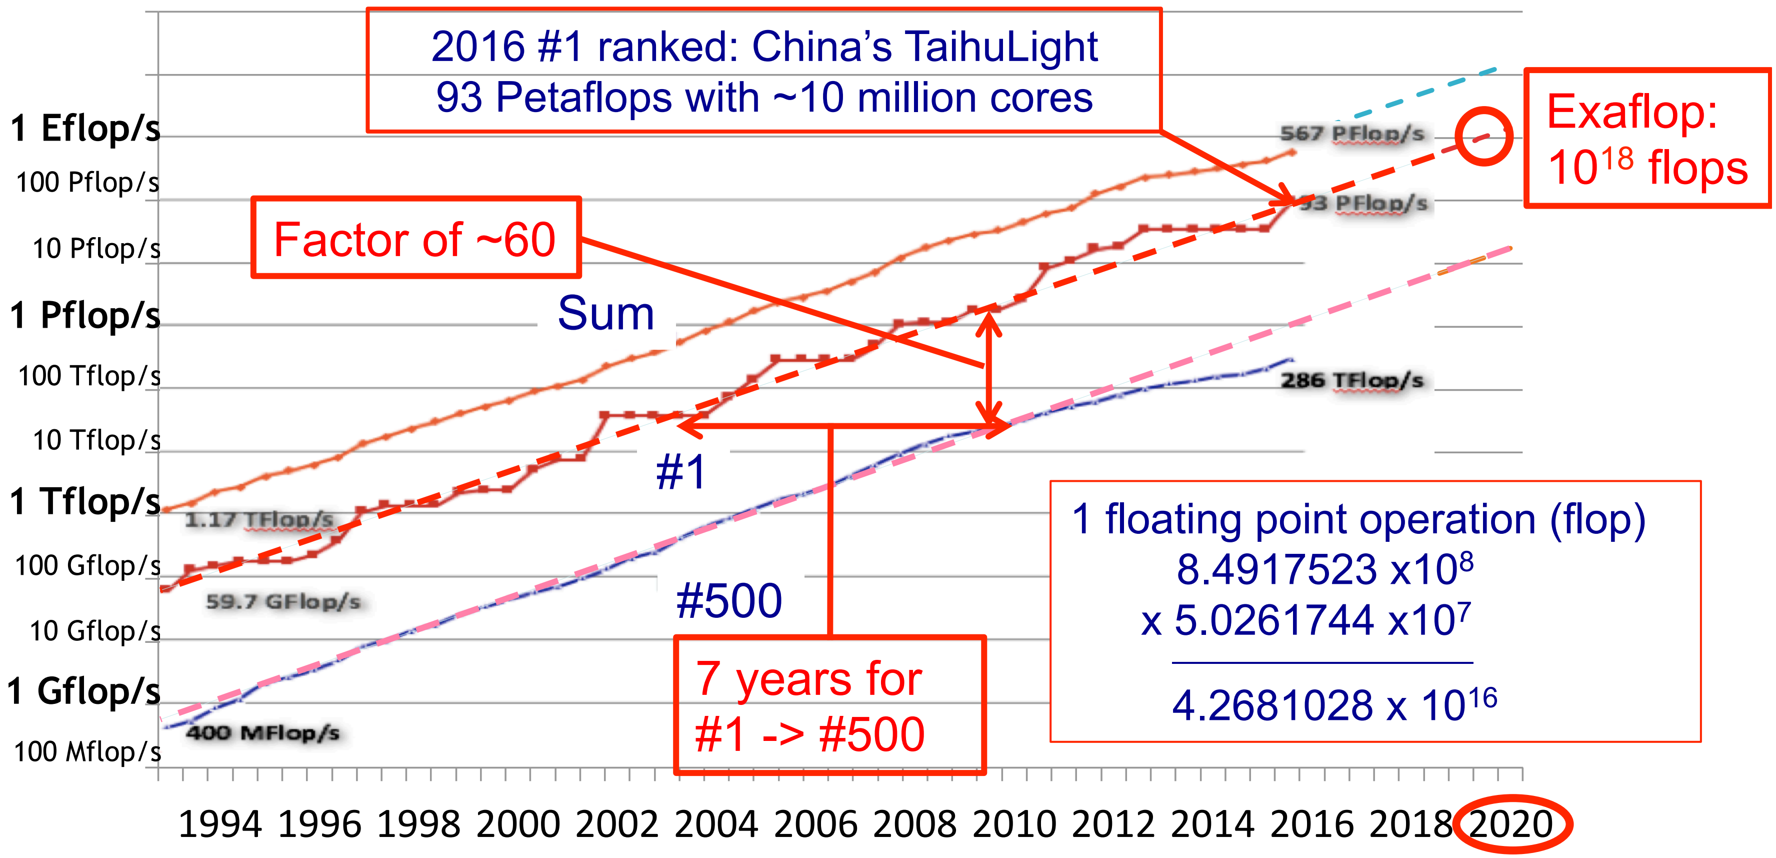

TOP500 SUPERCOMPUTER SITES

Adapted from Jack Dongarra's plot

HPC Research Representation at NTSE2016

NTSE2016 Participant	Top500	Ranking*
Junchen Pei, Peking University	Tianhe II	#2
James Vary, Iowa State University	Titan	#3
Noritaka Shimizu, Tokyo University	K-Super	#5
James Vary, Iowa State University	Mira	#6
Roman Skibinsky, Cracow University	Mira	#6
Vasily Kulikov, Moscow State University	Lomonosov II	#41
Andrey Shirokov, Moscow State University	Lomonosov II	#41
Andrey Shirokov, Moscow State University	Edison	#49
Youngman Kim, RISP	Edison	#49
Joel Lynn, Darmstadt University	Edison	#49
Xingbo Zhao, Institute of Modern Physics	Edison	#49
Charlotte Elster	Ruby & Owens	?
Youngman Kim, RISP	Tachyon II	?

NTSE2016 Participants applying for supercomputing		
Carlo Barbieri	Mira	#6
Luigi Coraggio	Fermi	#45

*June 2016 Top 500

US-DOE (ASCR) Computing Upgrades At a Glance

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Name Planned Installation	**Edison**	**TITAN**	**MIRA**	**Cori** **2016**	**Summit** **2017-2018**	**Theta** **2016**	**Aurora** **2018-2019**
System peak (PF)	3 – 30 PF			5 – 200 PF			
Peak Power (MW)	2 – 5 MW			2 – 14 MW			
Total system memory	0.3 – 0.8 PB			1 – 7 PB			
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~4,600 nodes	>3,200 nodes	>50,000 nodes
System Interconnect	Aries	Gemini	5D Torus	Aries	Dual Rail EDR-IB	Aries	2nd Generation Intel Omni-Path Architecture
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

Cray XK6 compute node

XK6 Compute Node Characteristics
AMD Opteron 6200 “Interlagos” 16 core processor @ 2.2GHz
Tesla M2090 “Fermi” @ 665 GF with 6GB GDDR5 memory
Host Memory 32GB 1600 MHz DDR3
Gemini High Speed Interconnect
Upgradeable to NVIDIA’s next generation “Kepler” processor in 2012
Four compute nodes per XK6 blade. 24 blades per rack

Intel Xeon Phi Knights Landing (just released)

Slide from Avinash Sodani (Intel), presented at ISC 2016 IXPUG workshop

Knights Landing Overview

TILE

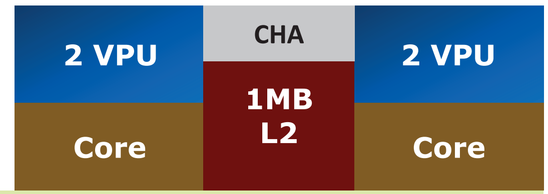

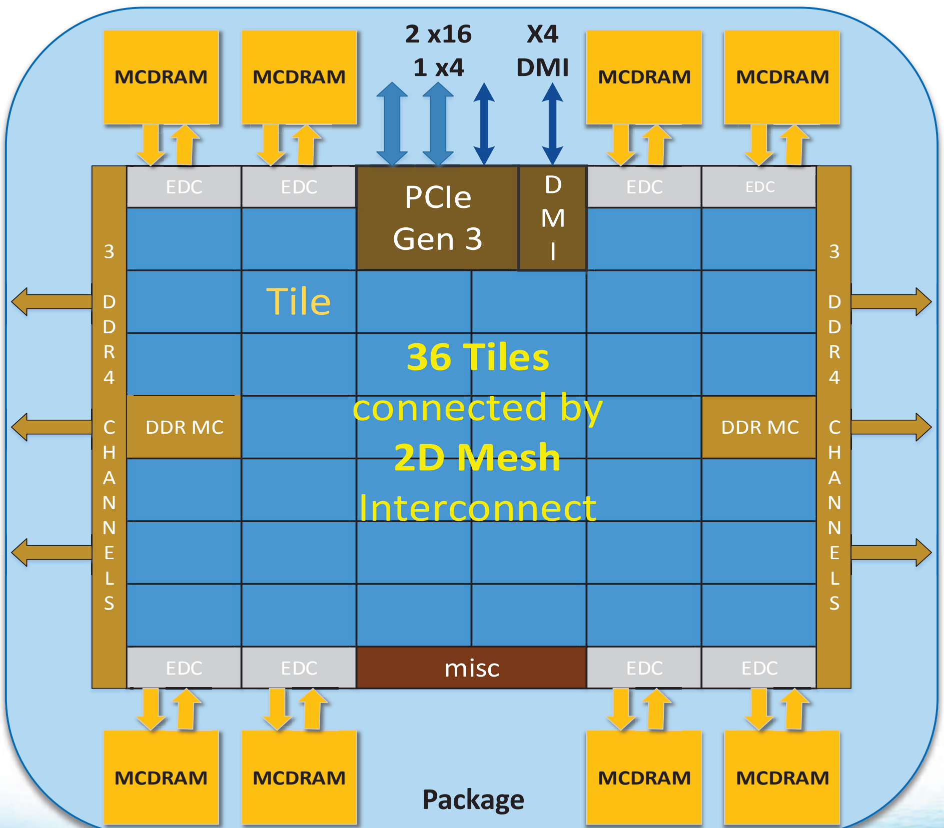

Chip: 36 Tiles interconnected by **2D Mesh**

Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW

DDR4: 6 channels @ 2400 up to 384GB

IO: 36 lanes PCIe Gen3. 4 lanes of DMI for chipset

Node: 1-Socket only

Fabric: Omni-Path on-package (not shown)

Vector[1]: up to 2 TF/s Linpack/DGEMM; 4.6 TF/s SGEMM

Streams Triad[1]: MCDRAM up to 490 GB/s; DDR4 90 GB/s

Scalar[2]: Up to ~3x over current Intel® Xeon Phi™ co-processor 7120 ("Knights Corner")

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit http://www.intel.com/performance. Configurations:

1. Intel Xeon Phi processor 7250 (16GB, 1.4 GHz, 68-cores) running LINPACK (score 2000 GFLOPS), DGEMM (score 2070 GFLOPS), SGEMM (4605 GFLOPS), STREAM (DDR4 = 90 GB/s and MCDRAM = 490 GB/s), 96 GB DDR4-2133 memory, BIOS R00.RC085, Cluster Mode = Quad, MCDRAM Flat or Cache, RHEL* 7.0, MPSP 1.2.2, Intel MKL 11.3.2, Intel MPI 5.1.2, DGEMM 20K x 20K, LINPACK 100K x 100K size
2. Intel estimates based on <specint-like workloads> comparing configuration 1 to Intel Xeon Phi co-processor 7120A hosted on 2x Intel Xeon processor E5-2697 v3.

Omni-path not shown

Node Capability Includes Memory Hierarchy

- Model of application memory use has changed
- What is the ***hottest*** memory? Coldest?

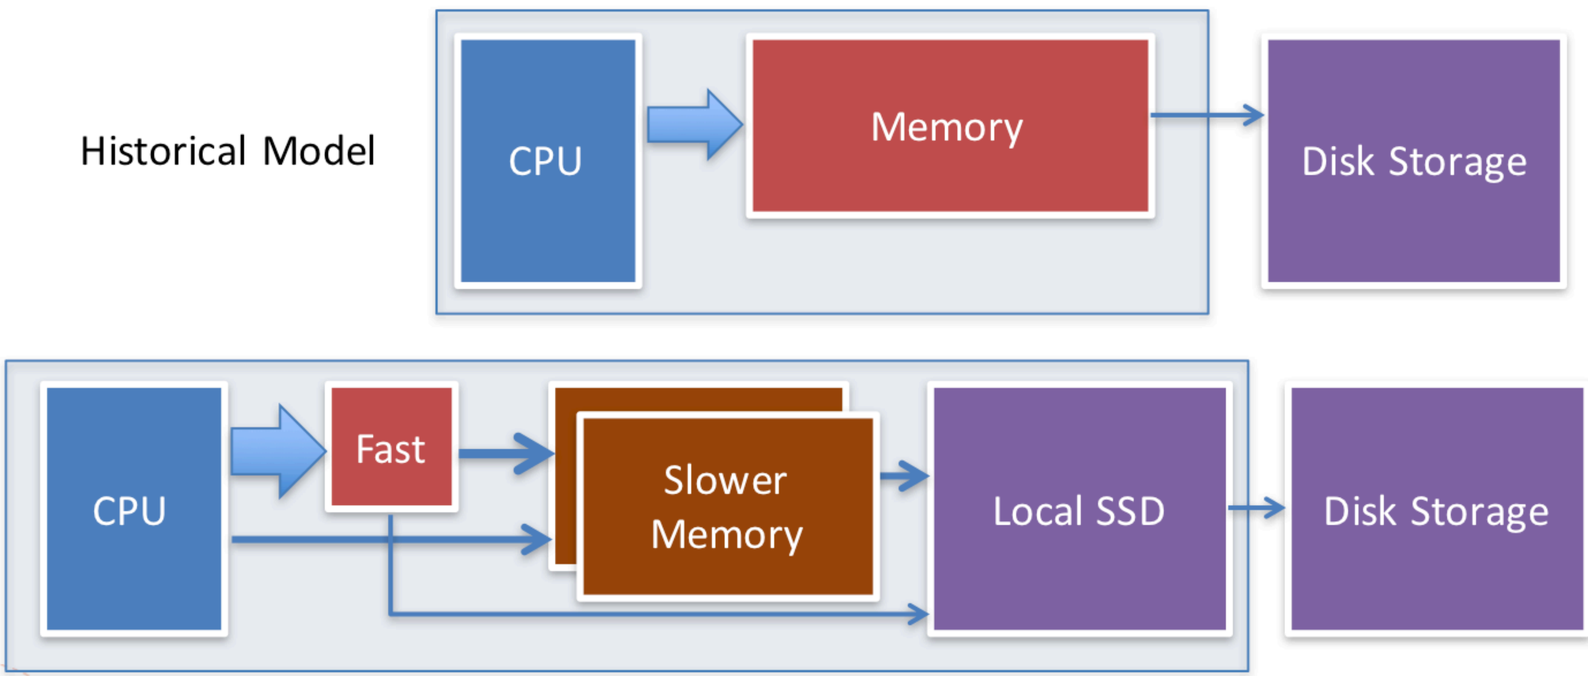

New model for many core or accelerated

Let's take a look at what is running on Edison now

NOW COMPUTING ON EDISON AT NERSC (http://www.nersc.gov/)

A small sample of massively parallel scientific computing jobs running right now at NERSC.

PROJECT	MACHINE	CPU CORES	CPU CORE HOURS USED
Chombo-Crunch: Advanced Simulation of Subsurface Flow and Reactive Transport Processes Associated with Carbon Sequestration PI: David Trebotich, Lawrence Berkeley National Laboratory	Edison	57,024	109,007.1
Controlling the Self-assembly of Alzheimer's-related Peptides: Efficient Sampling of Influencing Factors PI: Phillip L. Geissler, University of California Berkeley	Edison	5,184	113,657.4
Simulation of laser-plasma particle accelerators PI: Jean-Luc Vay, Lawrence Berkeley National Laboratory	Edison	4,800	104,415.7

NERSC Help Line: 1-800-666-3772

DID YOU KNOW?

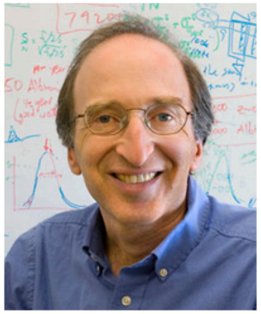

Saul Perlmutter—a professor of physics at UC Berkeley and a faculty senior scientist at Berkeley Lab—was awarded the 2011 Nobel Prize in Physics for his 1998 discovery that the universe is expanding at an accelerating rate. He confirmed his observations by running thousands of simulations at NERSC, and his research team is believed to have been the first to use supercomputers to analyze and validate observational data in cosmology.

Realities for Basic Science Supercomputers

Supercomputers for Basic Science Require Advance Planning

- What are the basic science “drivers”?
- How will simulations work on evolving architectures?
- Personnel development (Basic and Applied Science)?

Scientific/Societal/Political Efforts Intertwined in the Planning

- Scientific Community Consensus
- Funding Agency Support
- Government Support
- Popular Support

Requirements Gathering

LCFs and NERSC support a diverse user community

- Science benefits and impact of future systems are examined on an ongoing basis
- ASCR Computing Facility staff have been actively engaged in community assessments of future computational needs and solutions
- Computational science roadmaps are developed in collaboration with leading domain scientists
- Detailed performance analyses are conducted for applications to understand future architectural bottlenecks
- Analysis of INCITE, ALCC, ERCAP, Early Science and CAAR projects history and trends

Science Category	Represented Research Areas
Biology	Bioinformatics Biophysics Life Sciences Medical Science Neuroscience Proteomics Systems Biology
Chemistry	Chemistry Physical Chemistry
Computer Science	Computer Science
Earth Science	Climate Geosciences
Engineering	Aerodynamics Bioenergy Combustion Turbulence
Fusion	Fusion Energy Plasma Physics
Materials	Materials Science Nanoelectronics Nanomechanics Nanophotonics Nanoscience
Nuclear Energy	Nuclear Fission Nuclear Fuel Cycle
Physics	Accelerator Physics Astrophysics Atomic/Molecular Physics Condensed Matter Physics High Energy Physics Lattice Gauge Theory Nuclear Physics Solar/Space Physics

Argonne NATIONAL LABORATORY · BERKELEY LAB · OAK RIDGE National Laboratory

Slide from Katherine Riley, DOE

Why Exascale?

National Academy of Sciences study **NP2010:**

Nuclear Physics
Exploring the Heart of Matter

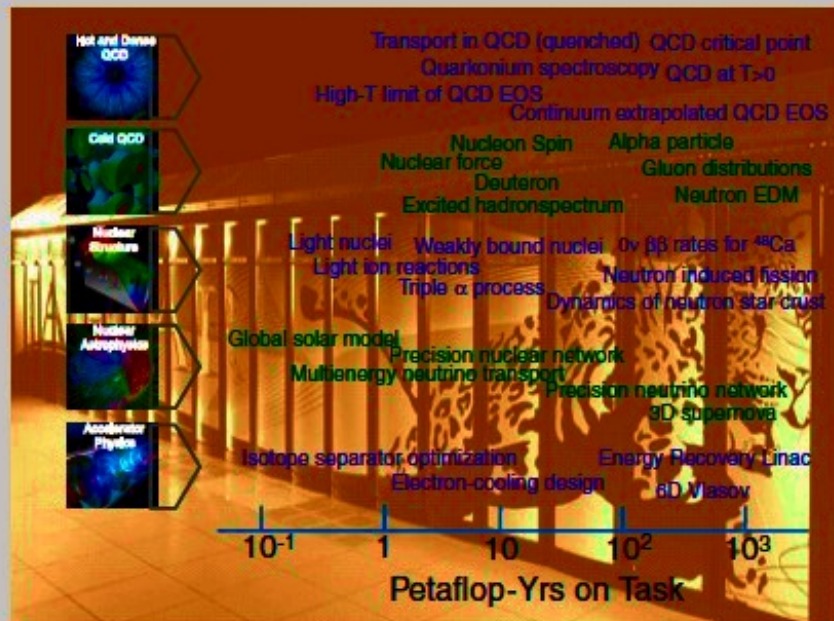

FIGURE 2.3.1 Estimates of the computational resources required to make breakthrough predictions in key areas of nuclear physics: hot and dense QCD, structure of hadrons, nuclear structure and reactions, nuclear astrophysics, and accelerator physics. SOURCE: Adapted from a figure by Martin Savage, University of Washington.

DOE National Laboratories Responses: 135
Academic Responses: 94
Industry Responses: 8
Foreign Response: 2
Others: 5

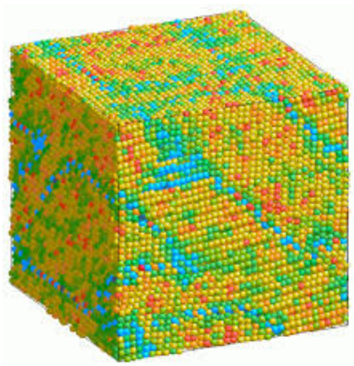

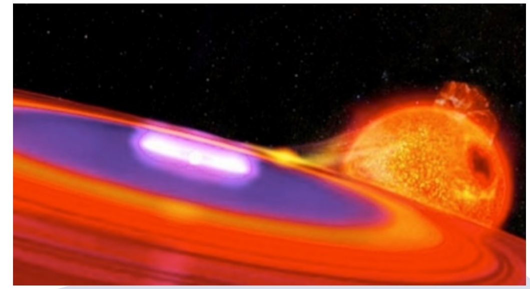

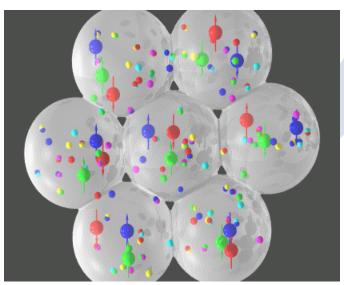

Nuclear, Particle & Plasma Physics

Chemical and Material Sciences

Biological Sciences & Public Health

Geoscience & Atmospheric Sciences

Astrophysics & Cosmology

Impacts Highlighted in Responses:

- Assuring National CyberSecurity
- Self Assembly based Nano-manufacturing
- Galaxy formation and extreme gravitational fields
- Optimizing the Power Grid
- Realistic Hypersonic flow for Flight Vehicles
- High Resolution Atmospheric & Climate Models
- N-by-N Comparison of All Patients in US
- Time-sensitive simulations of global scale epidemics

Slide from Barbara Helland, DOE

Terascale and Petascale simulation has facilitated steady scientific progress

Simulation is central to developing an understanding of planetary circulation response to radiative forcing changes

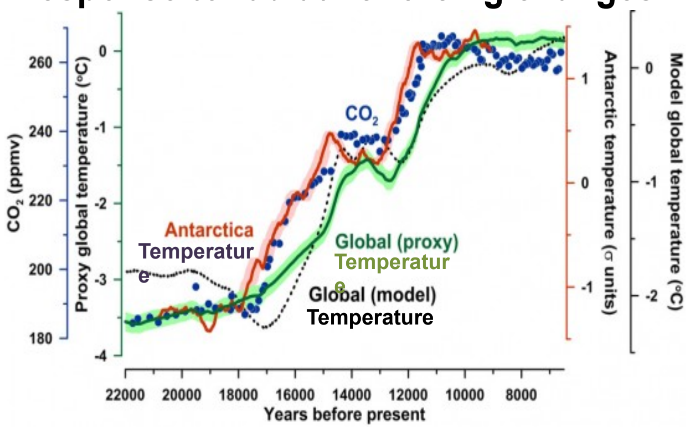

22,000 year simulation explains lag of CO_2 behind Antarctic temperature records over glacial cycles (Shakun et al., Nature)

Petascale capabilities critical to enabling treatment of eddy processes required to characterize regional climate

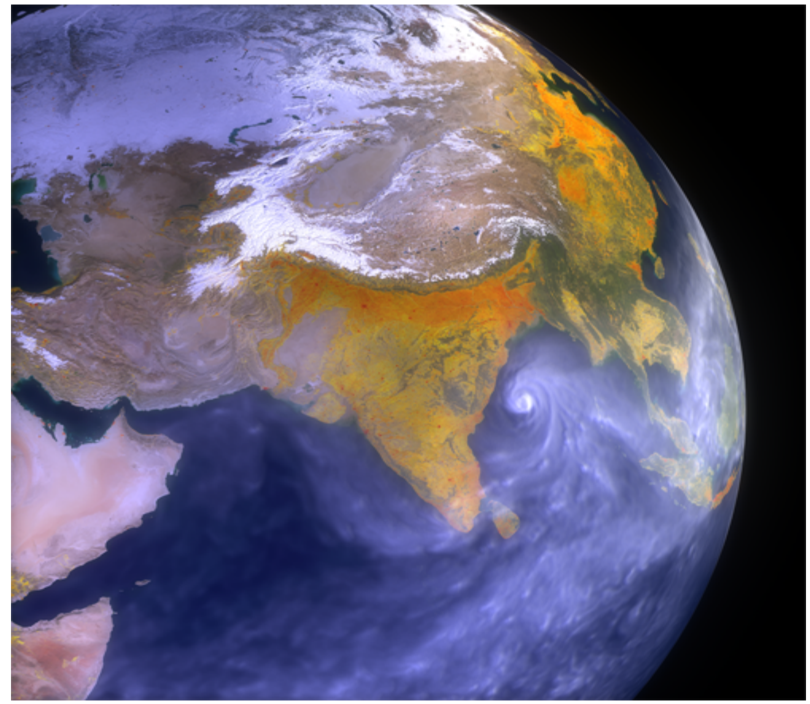

Petascale computing provides the capability to include motion scales in atmosphere and ocean required to address regional climate science questions in more complex frameworks. Snapshots from climate simulations illustrating simulation fidelity now within reach for climate applications

Slide from Barbara Helland, DOE

Exascale Will Advance Climate Science and Support Adaptation and Mitigation Studies

Science frontier:

- More accurate representations of cloud processes: fully resolve "eddy" motion scales, including cloud systems
- Better understanding and representation of Carbon-cycle feedbacks
- Better understanding and representation of aerosol and atmospheric **c**hemistry feedbacks
- Quantification of predictive skill as a function of motion and time scale

Predictive Science Outcomes:

- Climate variability on regional scales
  - Water cycle, water availability, food security
  - Frequency and intensity of extreme events
- Rate of Sea Level Rise
- Potential for abrupt change
- Interactions of societal responses with the natural system

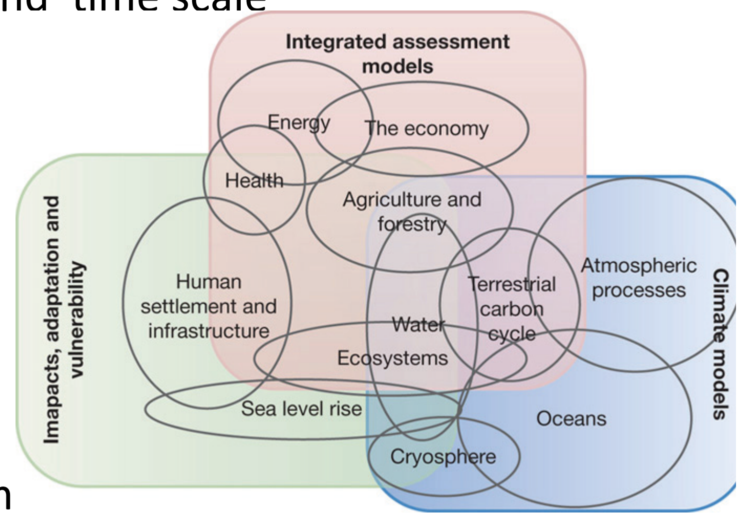

From 2015 BES Requirement Review Draft Report

BES Path to Exascale in Photosynthesis and Light Harvesting

- **Today:** Perform *ab initio* modeling of electronic phenomena of systems with hundreds of atoms.
- **2020: (petascale)** Perform *ab initio* modeling of electron transport and electron-phonon coupling of complex organic polymer systems of the order of tens of thousands of atoms. These simulations model individual components, ignoring interfaces between components that affect the performance of the devices
- **2025: (exascale)** Develop a model that accurately describes the electronic charge transport in the polymers and across electrical contacts, as well as the thermal transport across the warm and cold fluids through the ceramic layers and the organic polymers at appropriate levels of detail.

DOE "Big Data" Challenges
Volume, velocity, variety, and veracity

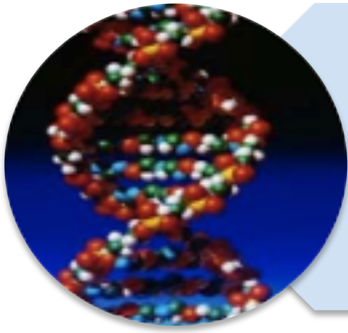

Biology
- Volume: Petabytes now; computation-limited
- Variety: multi-modal analysis on bioimages

Cosmology & Astronomy:
- Volume: 1000x increase every 15 years
- Variety: combine data sources for accuracy

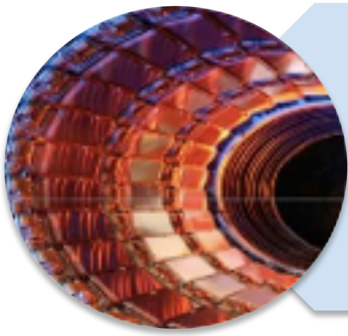

High Energy Physics
- Volume: 3-5x in 5 years
- Velocity: real-time filtering adapts to intended observation

Networks and Sensors:
- Variety: sensors for energy and network analysis
- Veracity: noisy data from carbon sensors, etc.

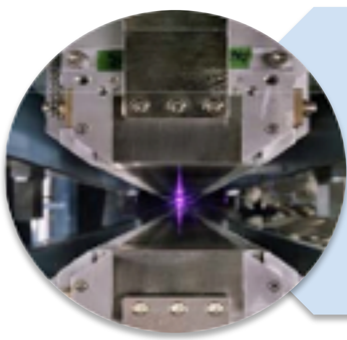

Light Sources
- Velocity: CCDs outpacing Moore's Law
- Veracity: noisy data for 3-D reconstruction

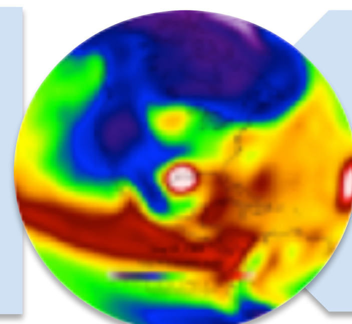

Climate
- Volume: Hundreds of exabytes by 2020
- Veracity: Reanalysis of 100-year-old sparse data

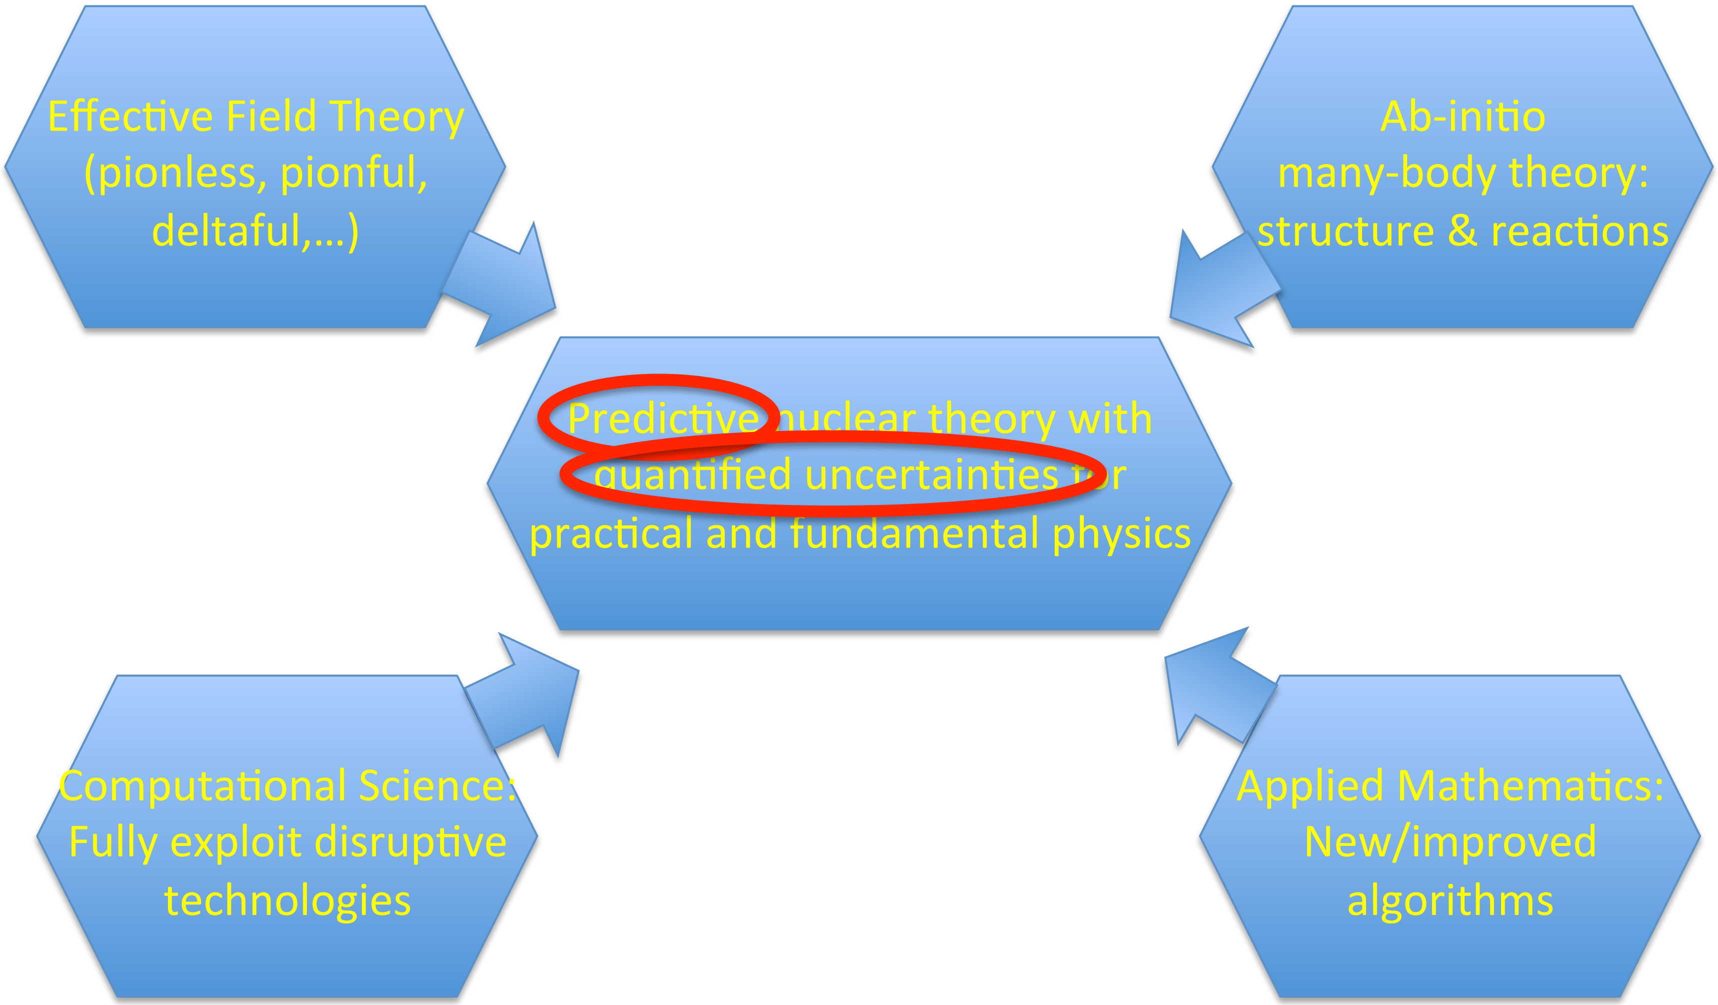
Effective Field Theory (pionless, pionful, deltaful,…)
Ab-initio many-body theory: structure & reactions
Predictive nuclear theory with quantified uncertainties for practical and fundamental physics
Computational Science: Fully exploit disruptive technologies
Applied Mathematics: New/improved algorithms

Cold QCD

Slide from Martin Savage and Joe Carlson

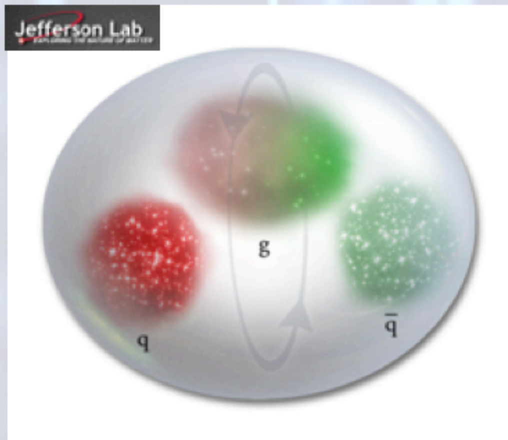

The Hadronic Spectrum and Exotics

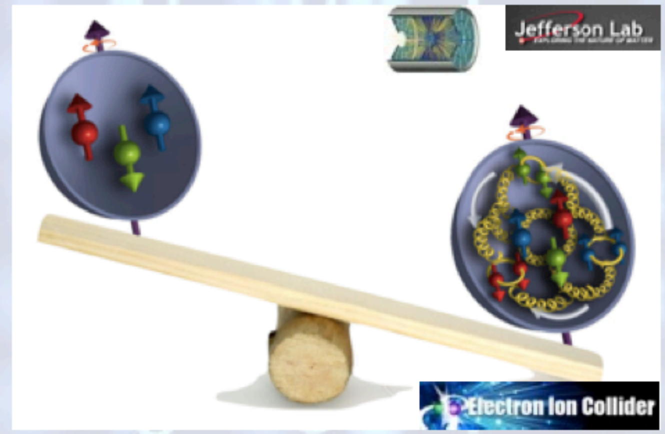

The Structure of the Proton

[DOE Topical Collaboration: TMD]

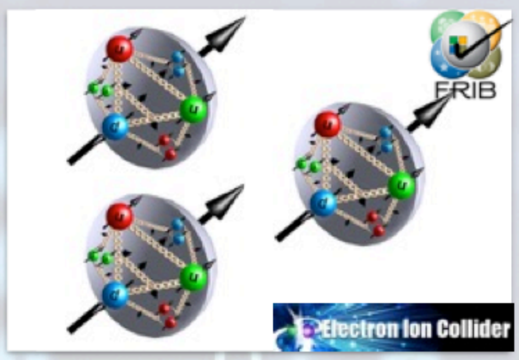

Nuclear Interactions and Forces

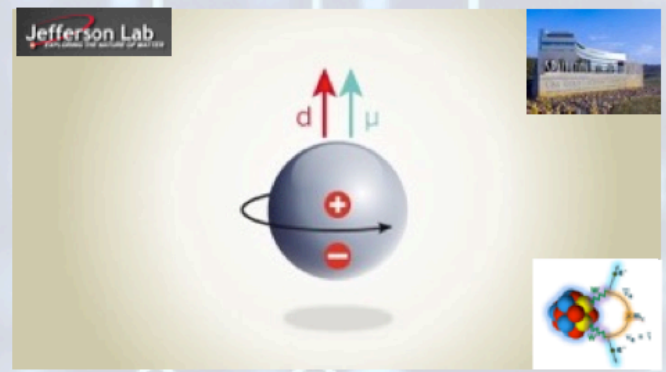

Fundamental Symmetries of Nature

Hot QCD

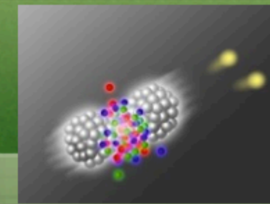

Slide from Martin Savage and Joe Carlson

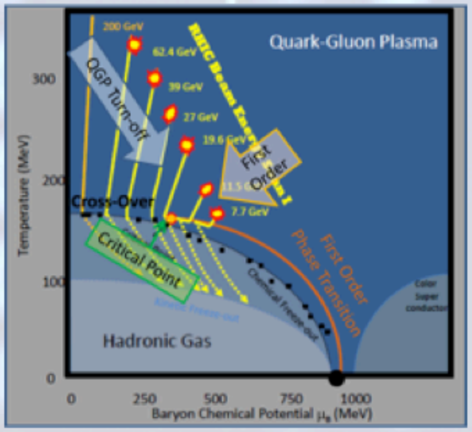

Properties of the QGP and Critical Point

[DOE Topical Collaboration: BES]

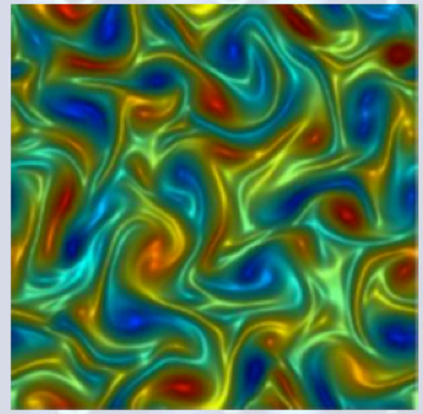

Equilibration

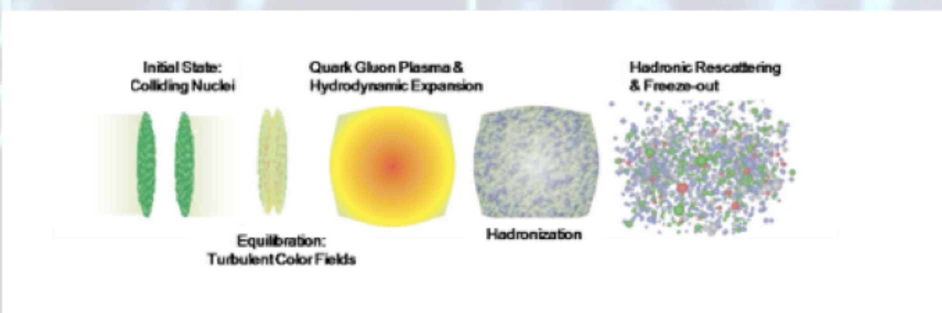

Modeling heavy-ion collisions

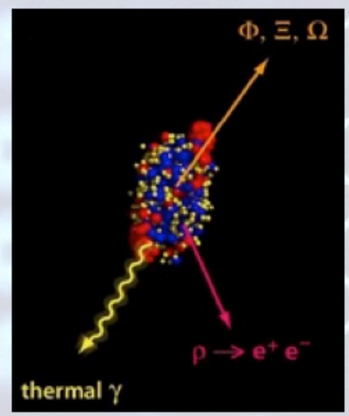

Spectral Functions and Hadron Properties

Nuclear Structure/Reactions

Slide from Martin Savage and Joe Carlson

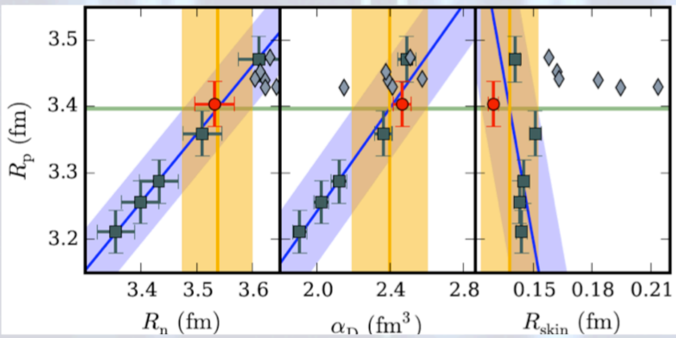

Nuclei from First Principles
Light Nucleus Reactions

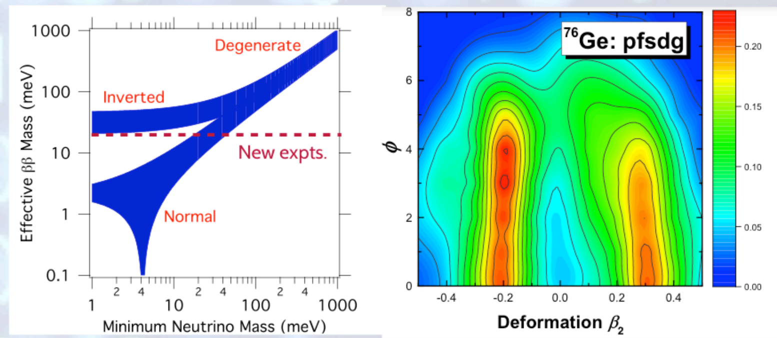

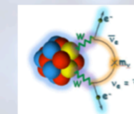

Electroweak Interactions

[DOE Topical Collaboration: Double Beta Decay and FS]

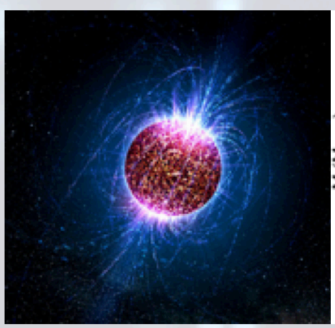

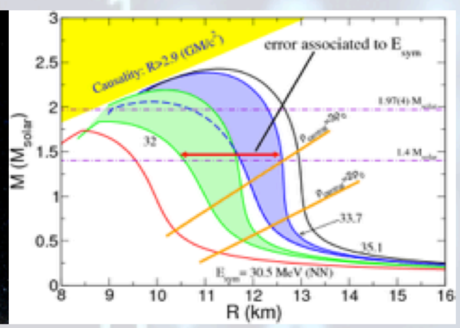

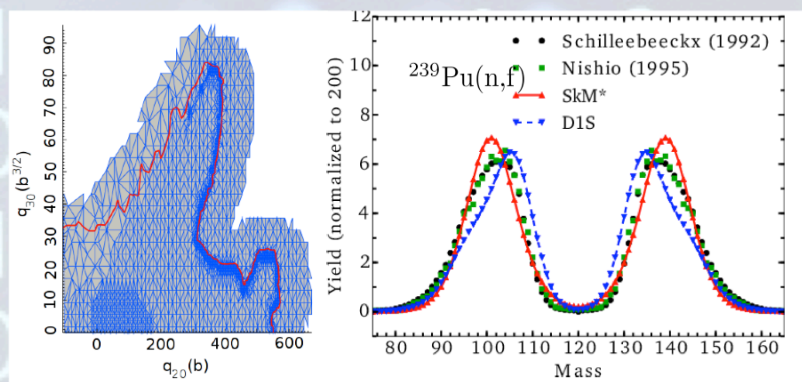

Dense Nuclear Matter

Density Functional Theory

Nuclear Astrophysics

Slide from Martin Savage and Joe Carlson

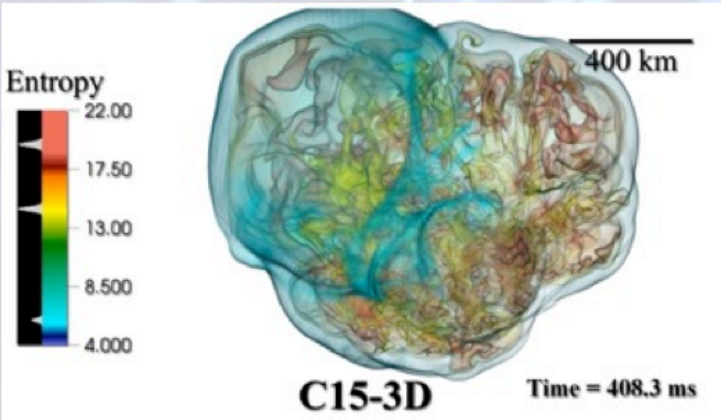

Core-Collapse Supernovae
Thermonuclear Transients

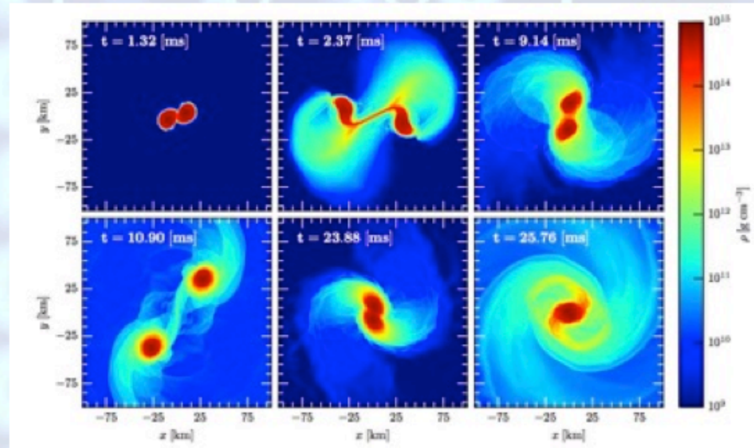

Neutron Star Mergers

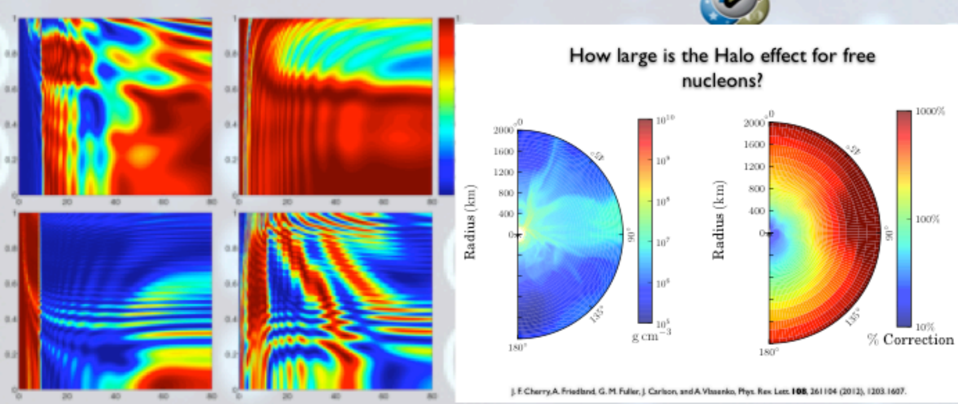

Neutrino Flavor Physics

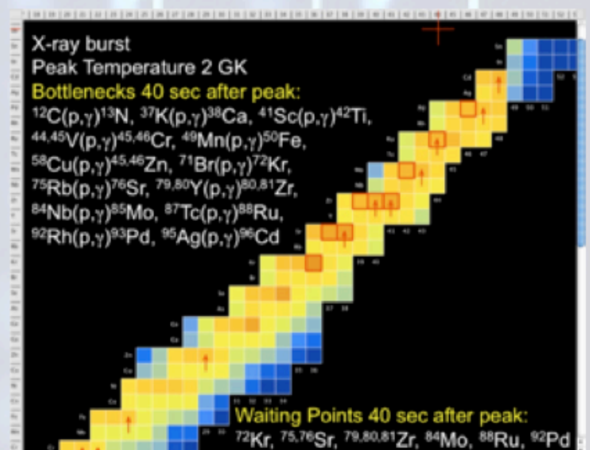

Explosive Nucleosynthesis

Experiment

Slide from Martin Savage and Joe Carlson

FRIB: Greta, Design Ops, DAQ

RHIC: STAR, SPhenix

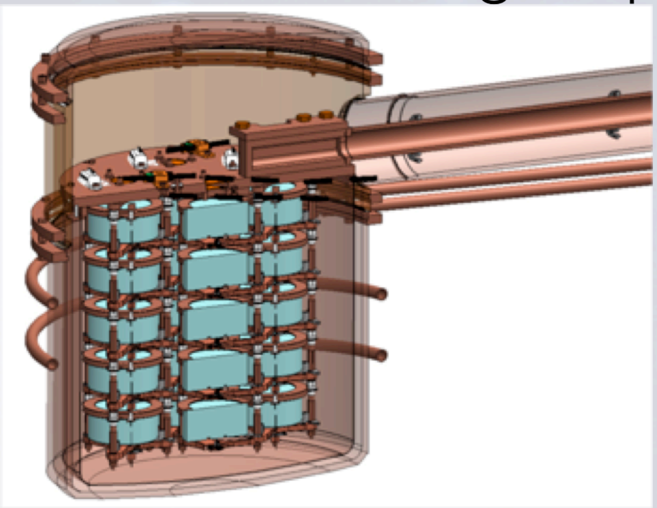

Double Beta Decay

Electron Ion Collider

JLab: electron cooling, beam-beam

Other facilities: ATLAS, FNPB (SNS), HIGS, University Facilities, …

Beyond Nuclear Physics

Slide from Martin Savage and Joe Carlson

Advanced LIGO
neutron star mass/radius

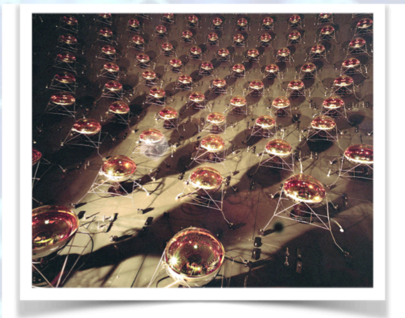

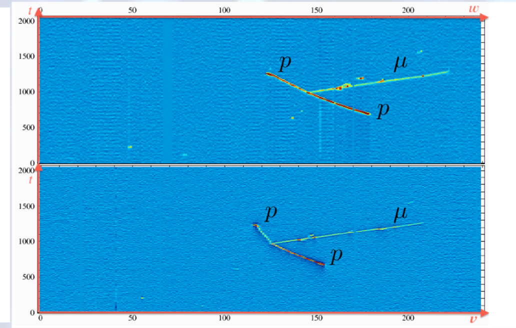

Neutrino Physics
ν-nucleus, CP, hierarchy

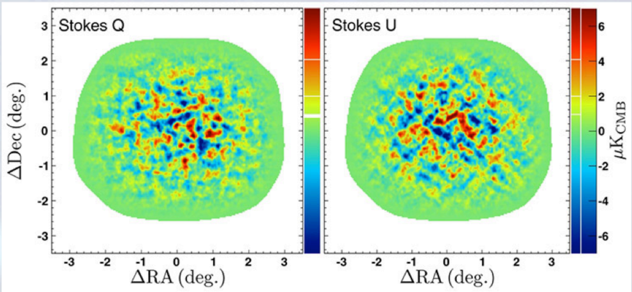

Neutrino masses and cosmic microwave background

HINP QCD phase diagram (calc. at m_B = 0) with physical light quark masses

"The QCD phase transition with ***physical-mass****, chiral quarks."* (hotQCD Collab.)

arXiv:1402.5175v1 (2014) **MAJOR** CPU time reqts.

No phase transition at physical m_q . (m=0)

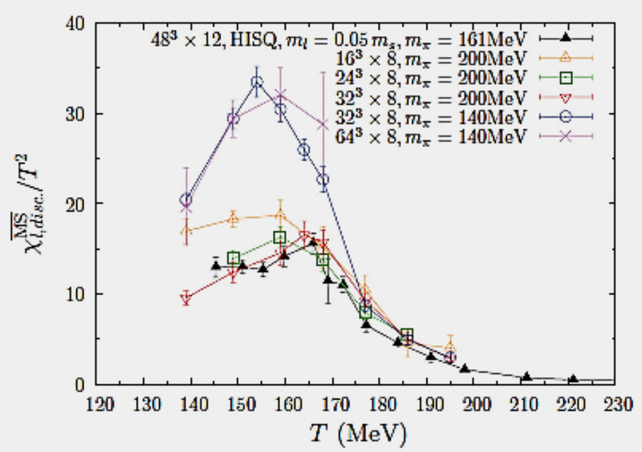

FIG. 1. The dependence of the disconnected chiral susceptibility on temperature for m_π = 135 and 200 MeV. The m_π = 135 MeV data shows a near 2× increase over that for m_π = 200 MeV. HISQ results for m_π = 161 MeV [7, 11] are also plotted.

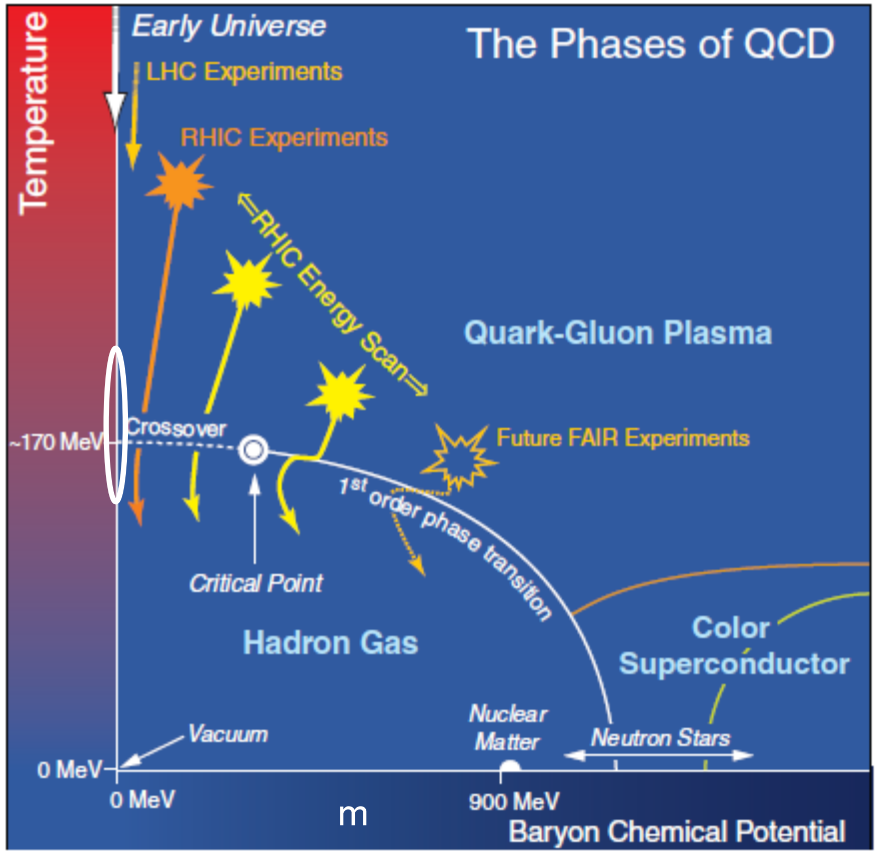

FIGURE 2.25 The phase diagram of QCD is shown as a function of baryon chemical potential (a measure of the matter to antimatter excess) and temperature. A prominent feature in this landscape is the location of the critical point, which indicates the end of the first-order phase transition line in this plane. SOURCE: DOE/NSF, Nuclear Science Advisory Committee, 2007, *The Frontiers of Nuclear Science: A Long Range Plan*.

Slide from Ted Barnes, DOE

U.S. DEPARTMENT OF ENERGY | Office of Science

<u>Now</u> working on extending this to finite chemical potential; m power series.

LENP

Two possible sites of r-process nucleosynthesis.

CCSN

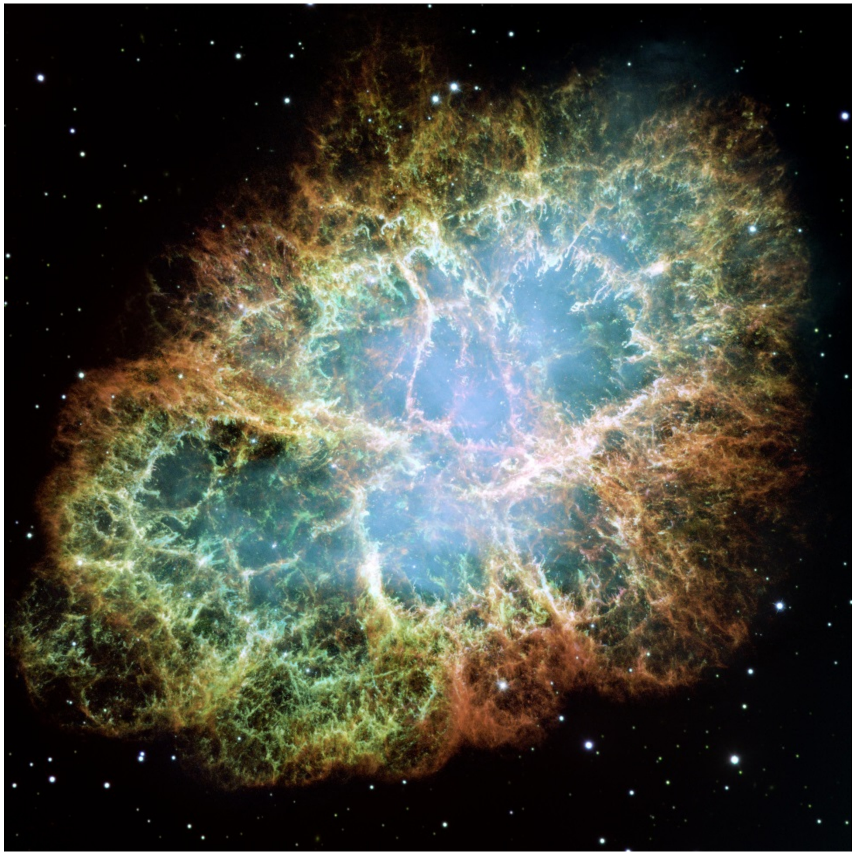

Crab Nebula supernova remnant (1054 AD).
Simulation of high density n physics?

Neutron star merger

Neutron star merger and the "chirp."

LENP

Three peaks; do we need both CCSNe *and* neutron star mergers?

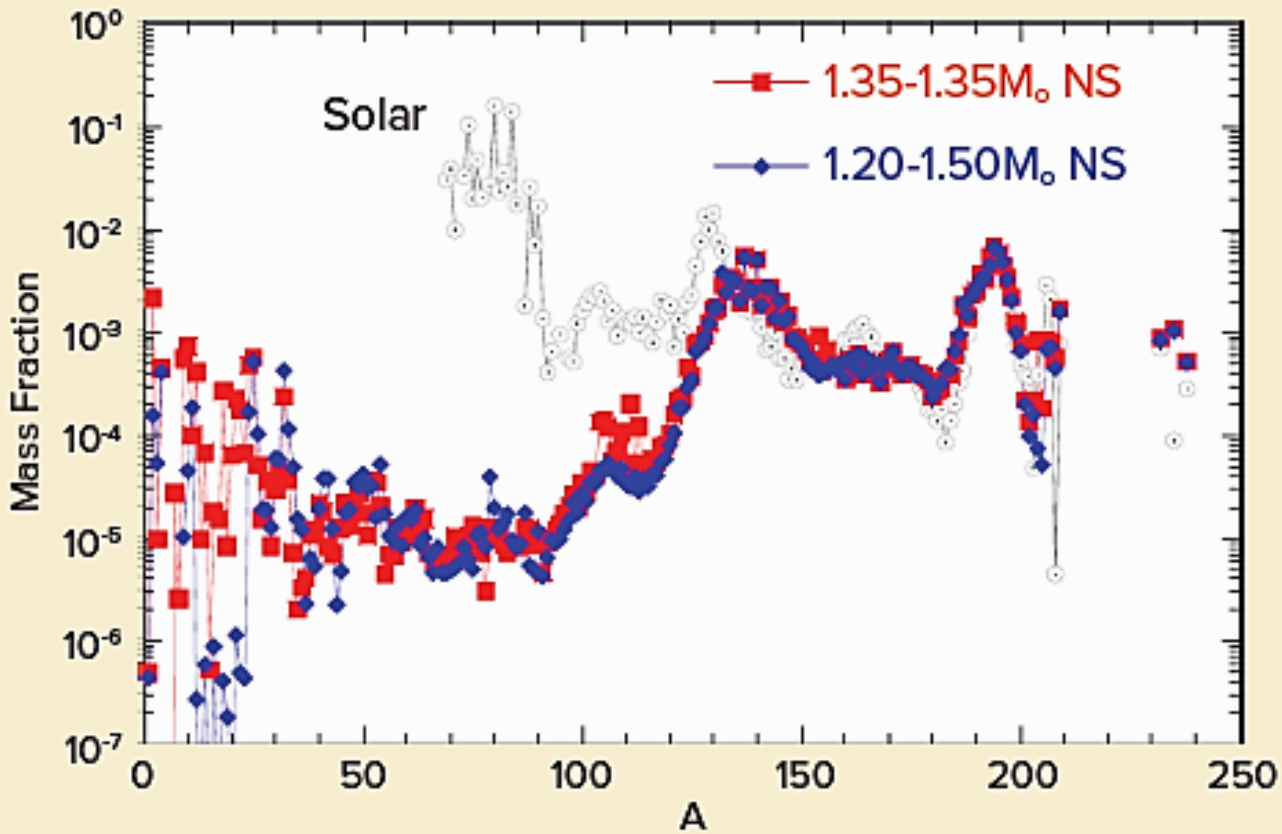

Nuclear abundance distributions as a function of atomic mass of the ejecta for two different combinations of neutron star mergers. The distributions are normalized to the solar r-process abundance distribution. Image credit: Stephane Goriely.

LENP

Allowed neutron star maximum masses for various EOSs.

Once you can calc. the nucl mat. EOS (NN, NNN, s), you will be able to determine allowed neutron star masses and compositions.

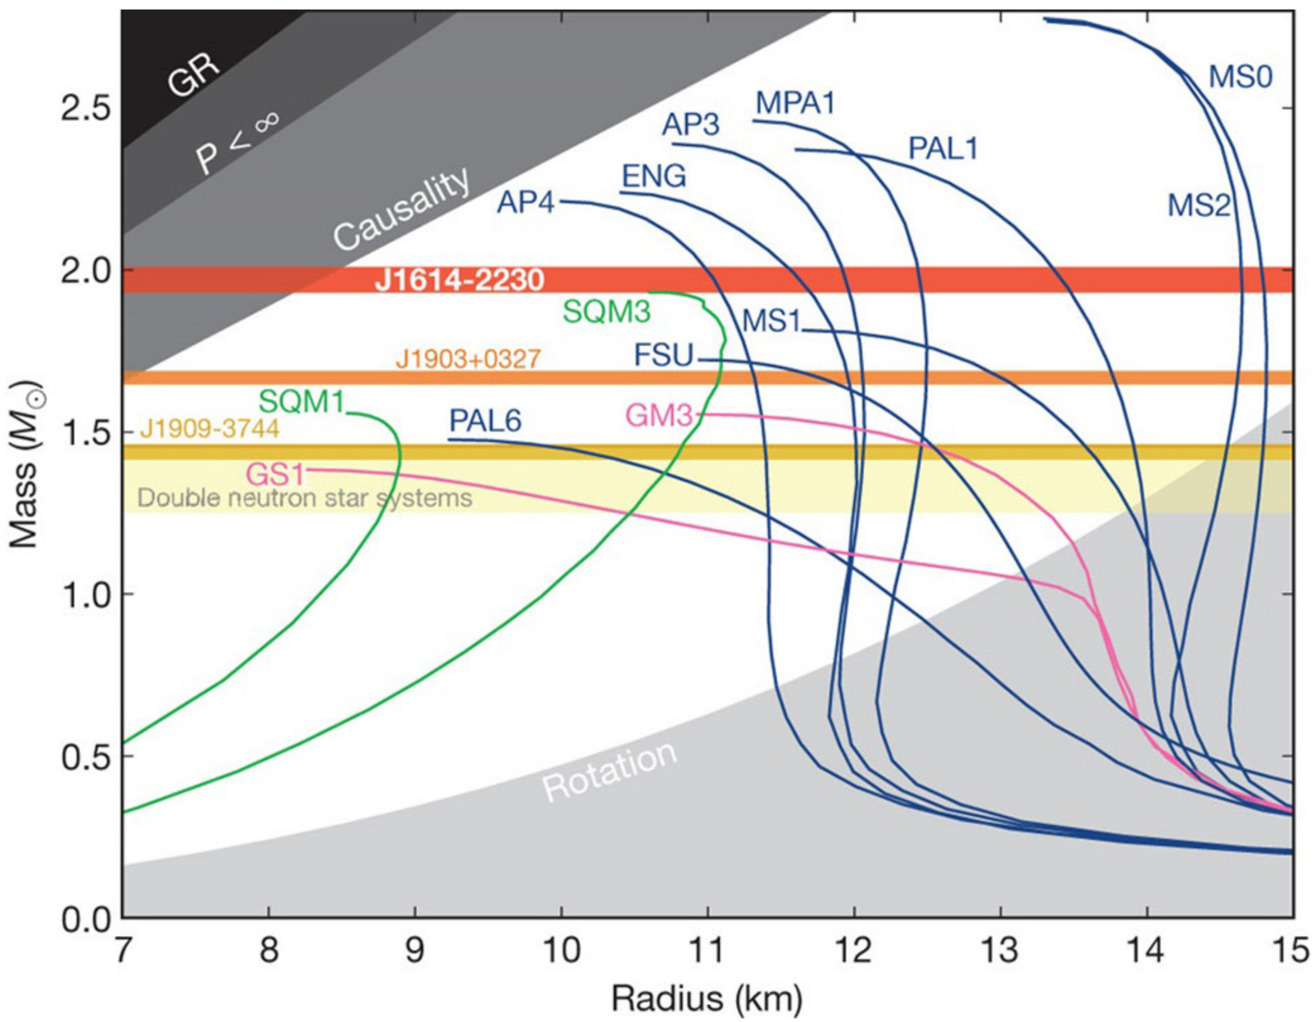

Fundamental questions of nuclear physics => discovery potential

- What controls nuclear saturation?
- How shell and collective models 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?

U.S. DEPARTMENT OF ENERGY | Office of Science

National Science Foundation
WHERE DISCOVERIES BEGIN

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
INCITE
Advancing America's Science and Industrial Competitiveness

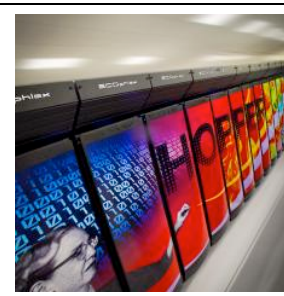

+ K-super.
+ Blue Waters
+ Lomonosov II
+ Tachyon-II

LENP

Nuclei, Nuclear Matter, Nuclear Astrophysics

- FRIB physics is at the core of nuclear science: "To understand, predict, and use"
- FRIB provides access to a vast unexplored terrain in the chart of nuclides

ca. 7000 nuclei exist (NUCLEI SciDAC-3 project)

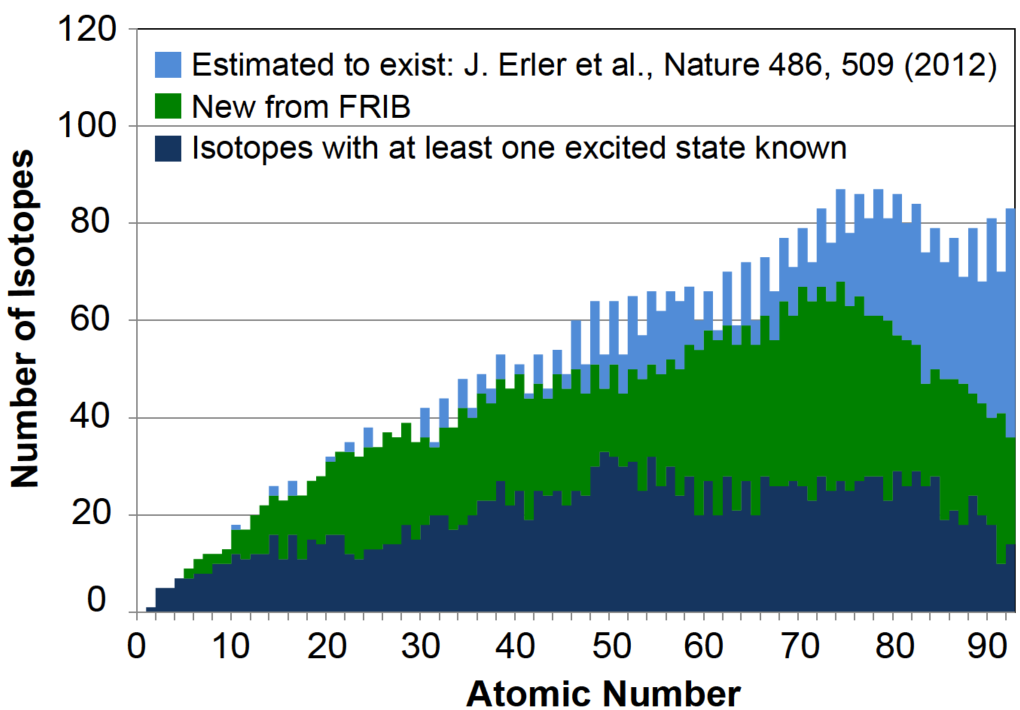

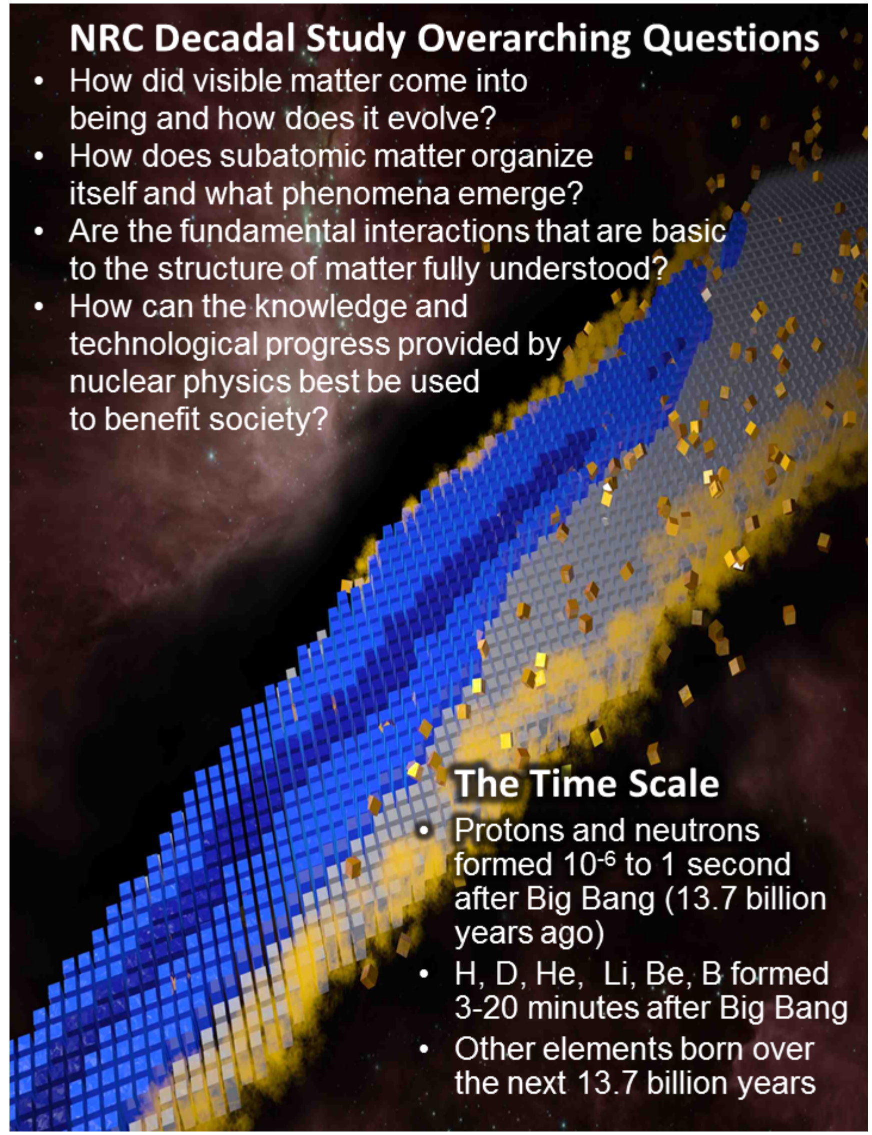

U.S. DEPARTMENT OF ENERGY | Office of Science

MANY LENP applications: NUCLEI SciDAC has over 200 pubs. in 3+ years. … hard to choose!

Slide from Ted Barnes, DOE

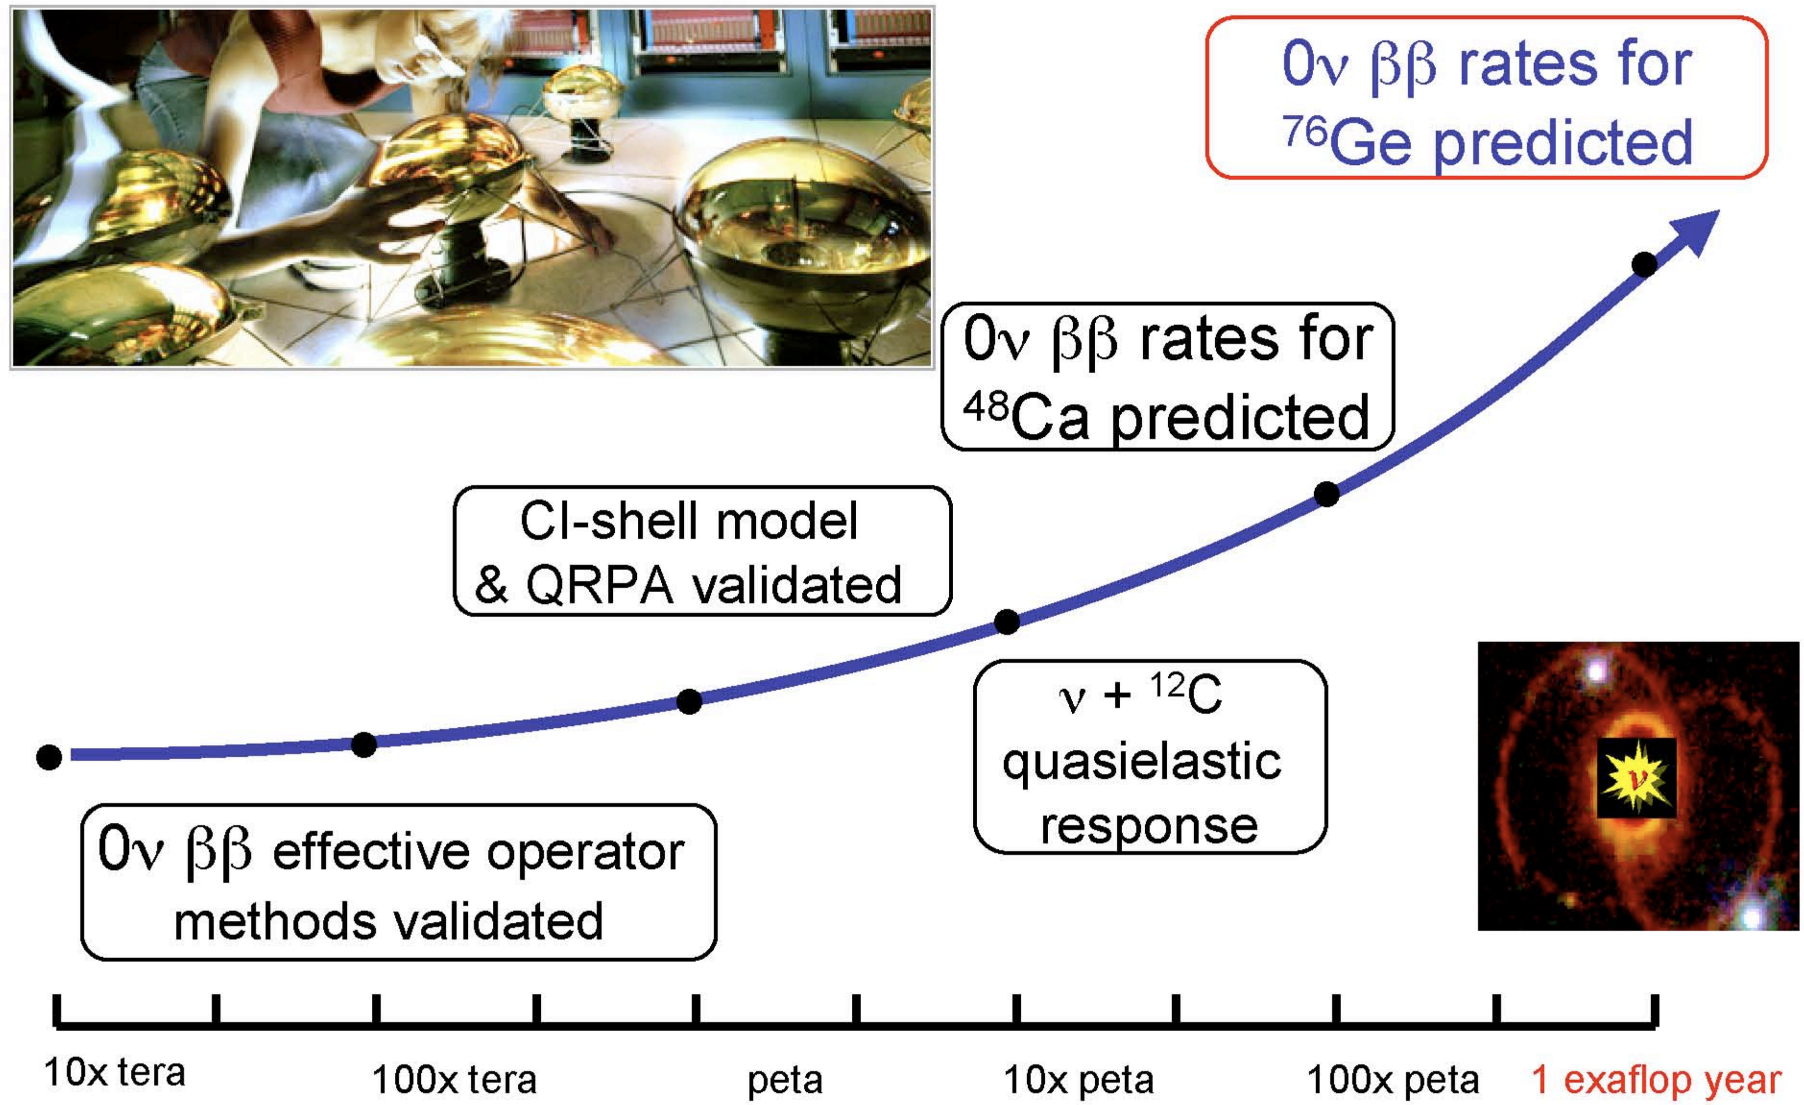

http://extremecomputing.labworks.org/nuclearphysics/report.stm

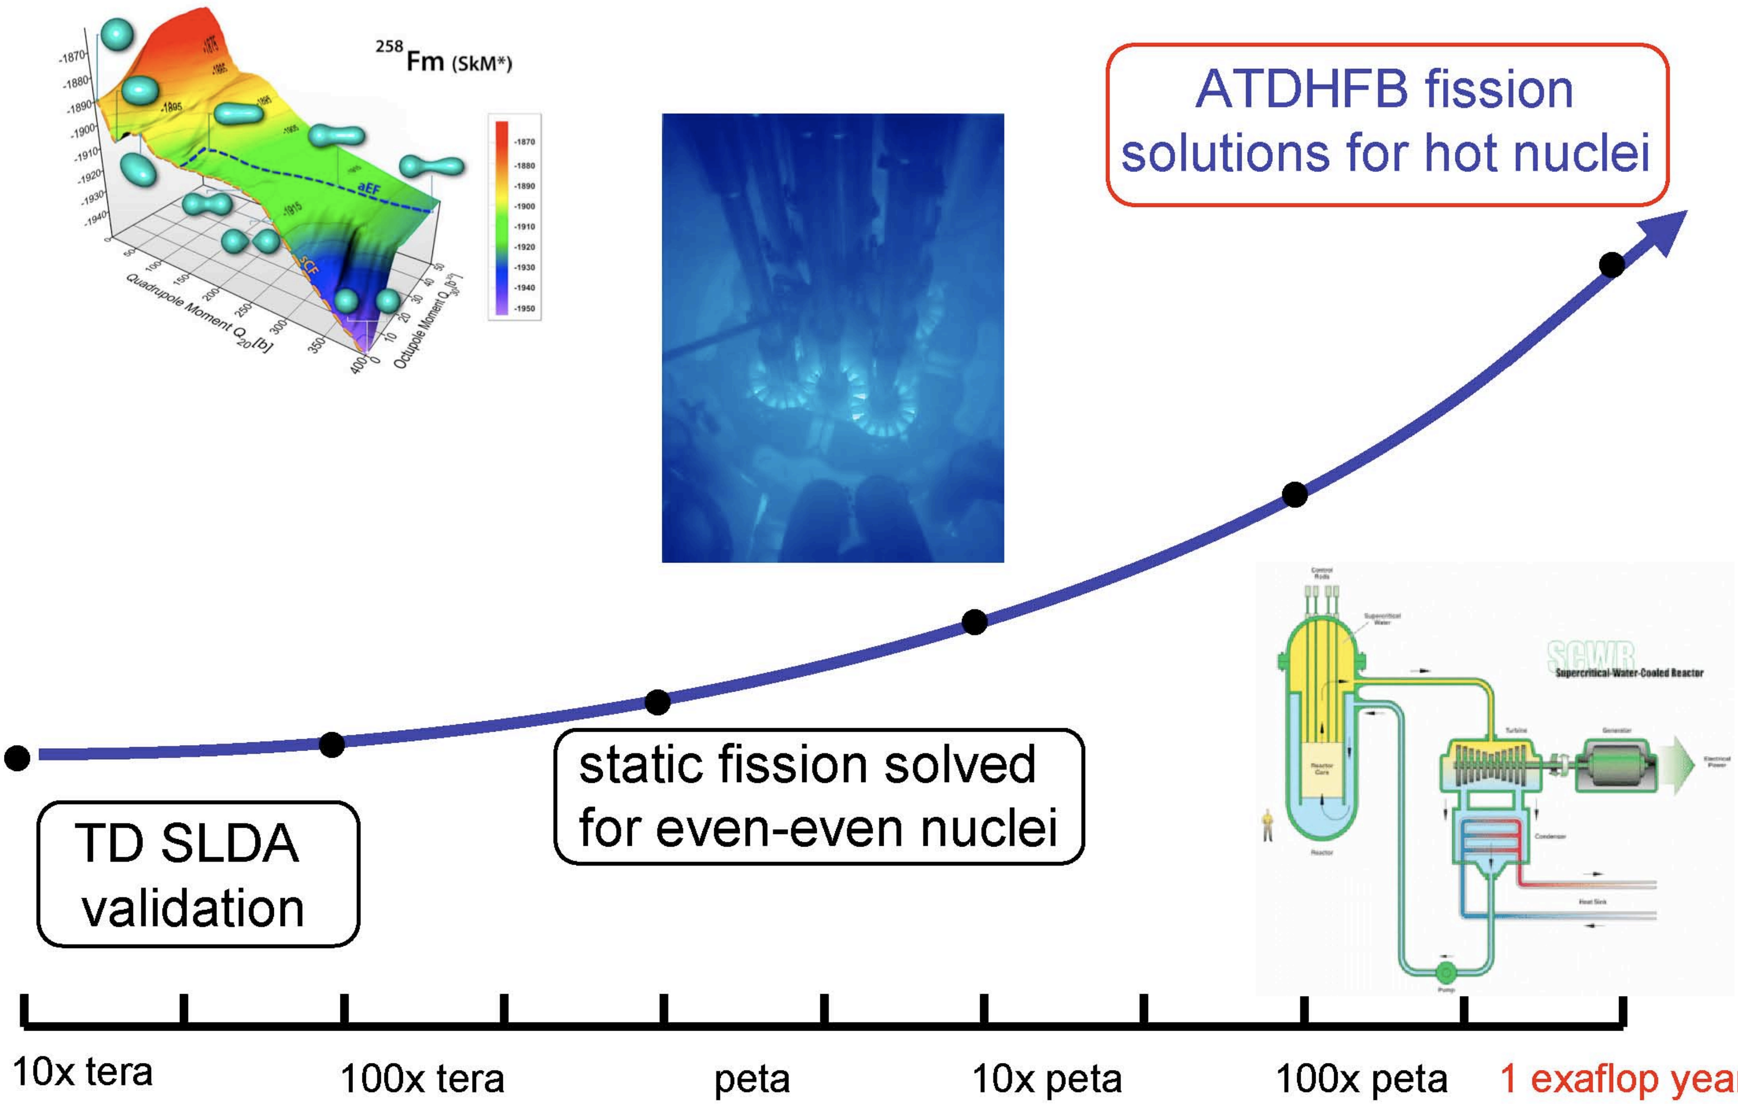

http://extremecomputing.labworks.org/nuclearphysics/report.stm

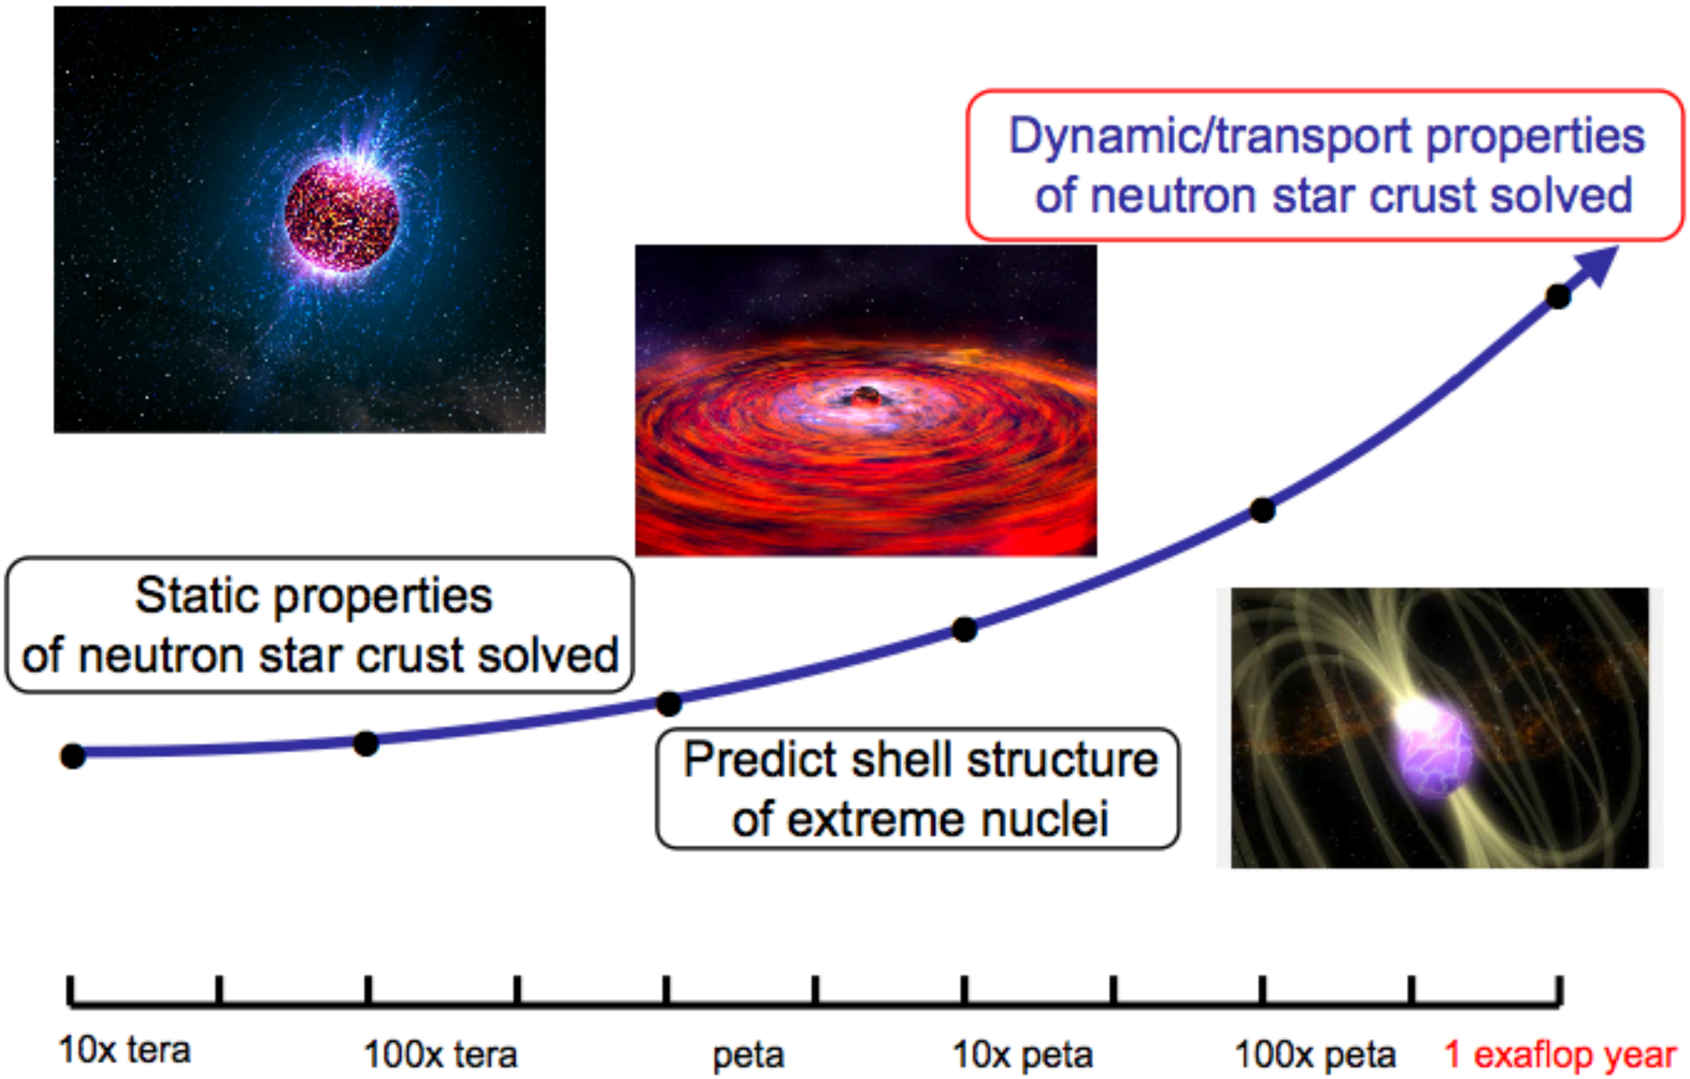

http://extremecomputing.labworks.org/nuclearphysics/report.stm

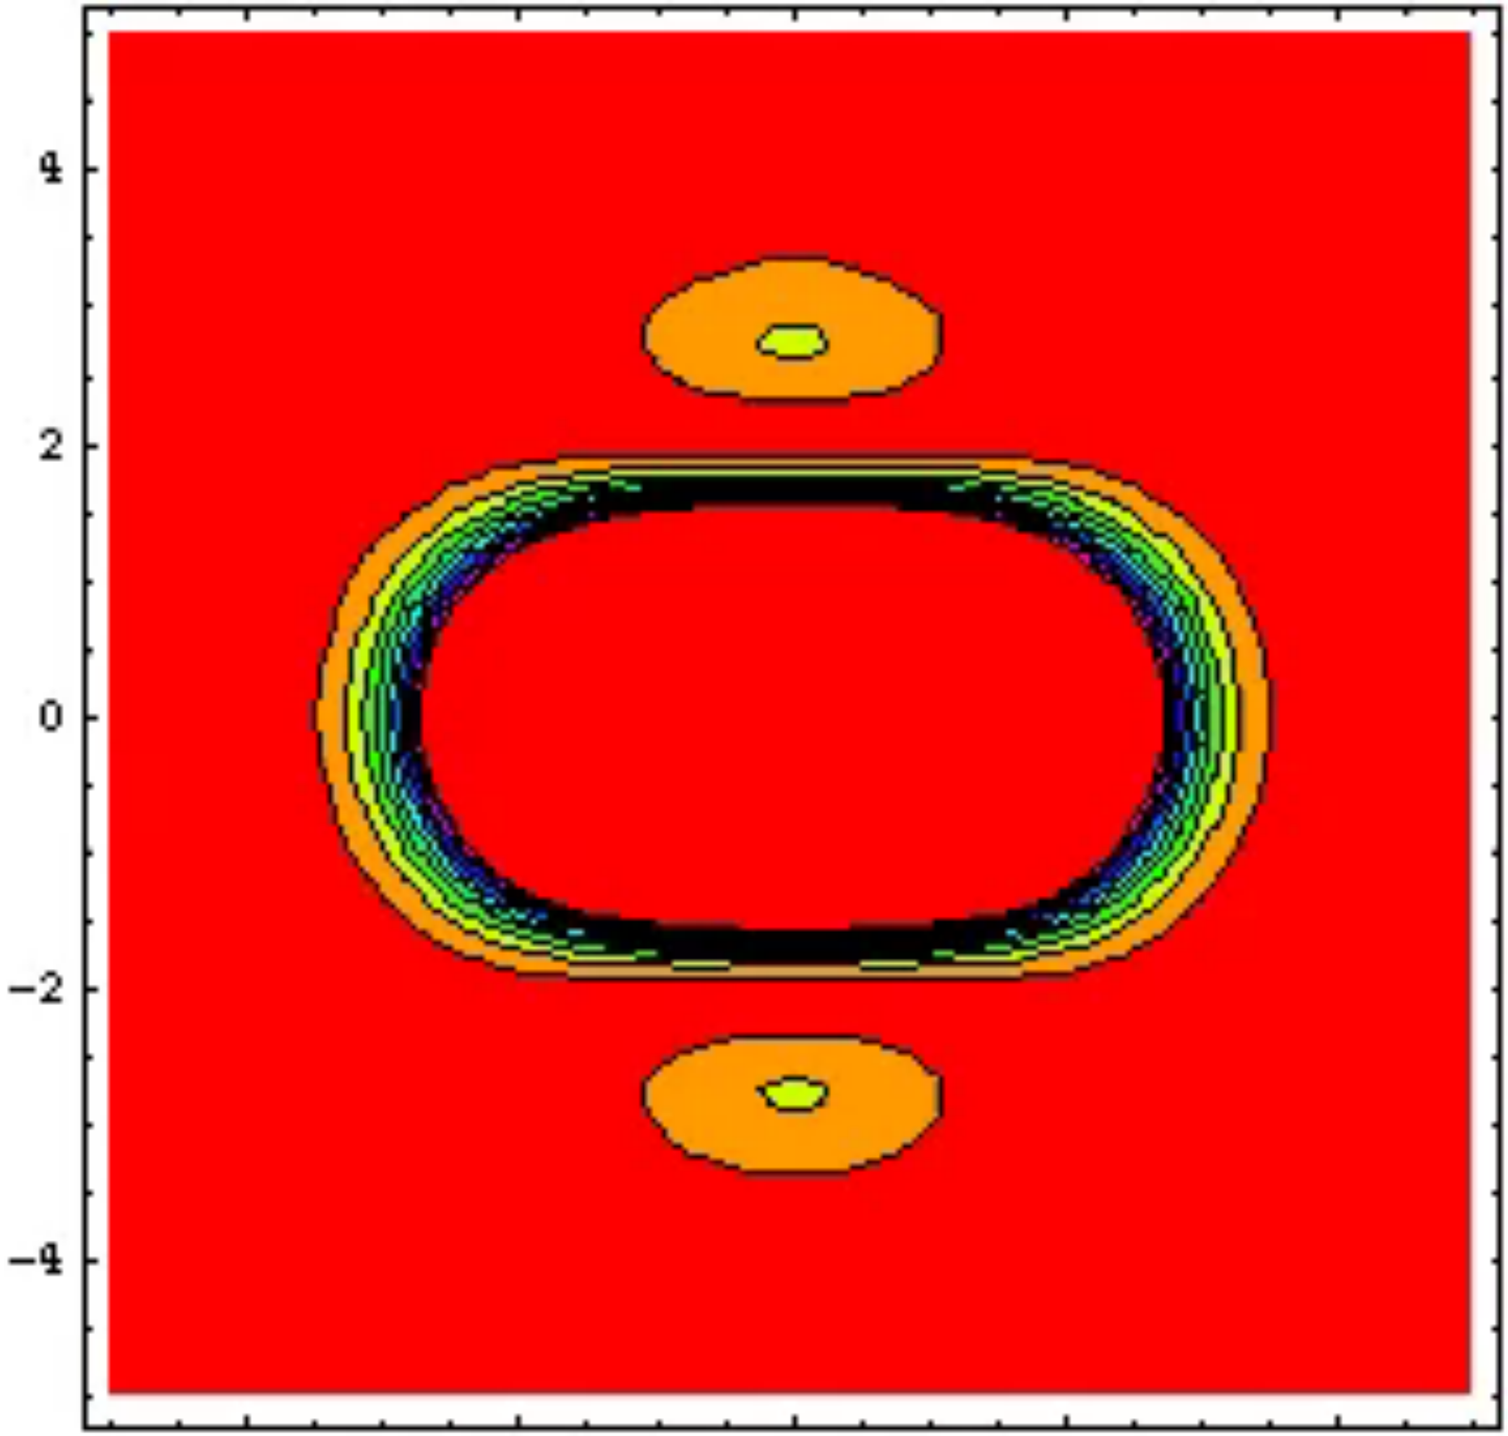
4
2
0
−2
−4

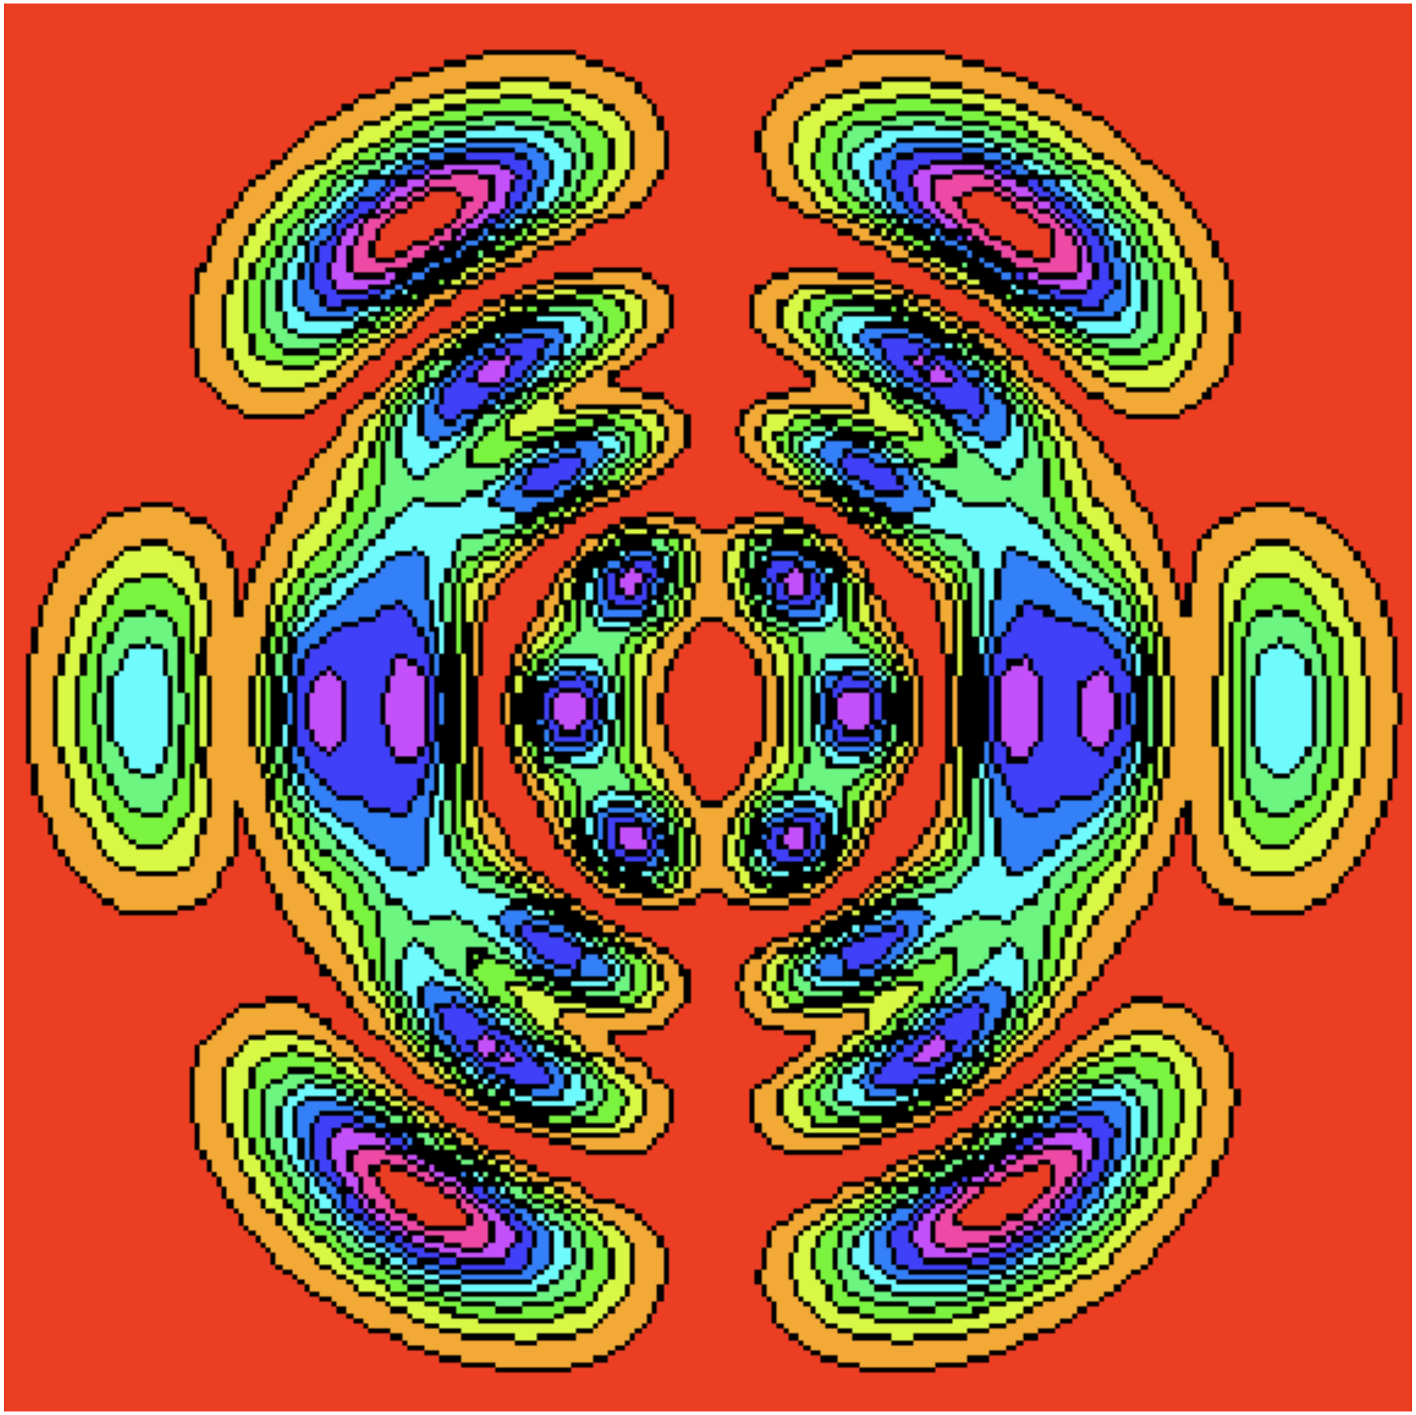

^{8}Li gs
J=2

Monopole   Quadrupole   Hexadecapole

Neutrons

Protons

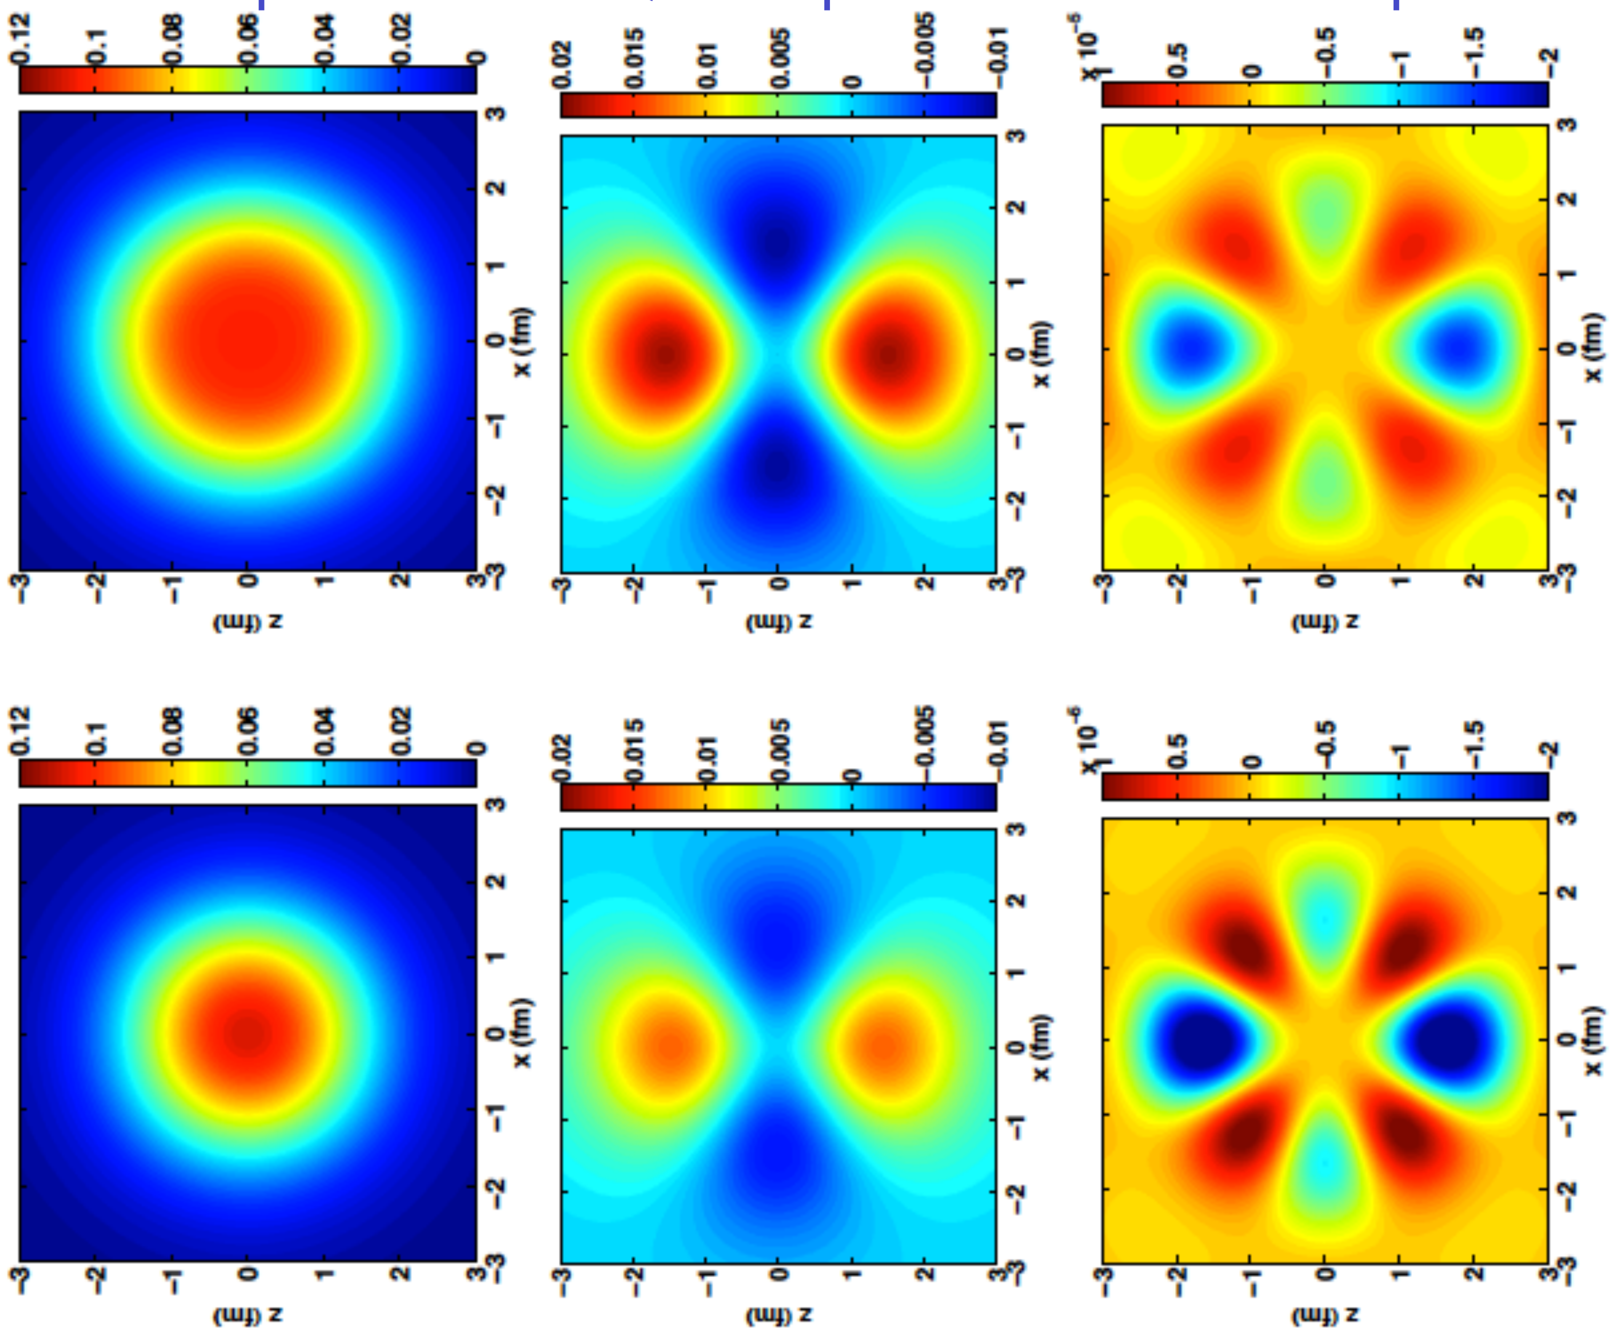

C. Cockrell, J.P. Vary, P. Maris, Phys. Rev. C86, 034325 (2012)
C. Cockrell, PhD, Iowa State University

One-body density of 9Be ground state $(\frac{3}{2}^-, \frac{1}{2})$

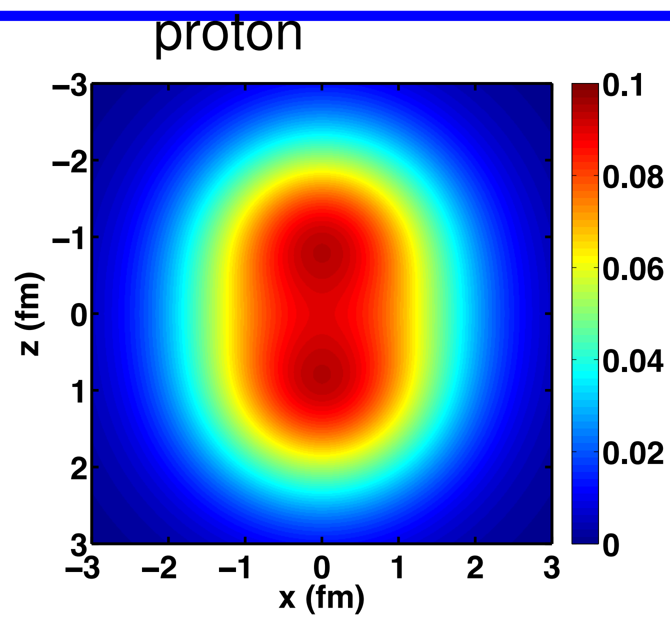

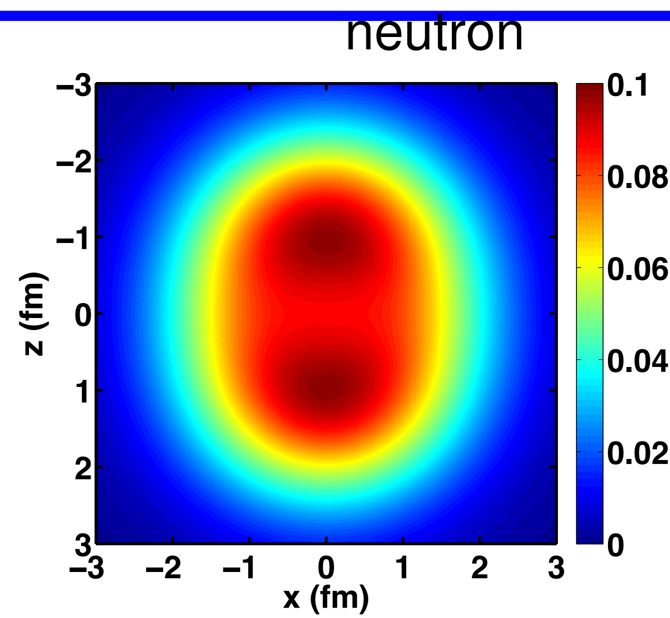

and their difference

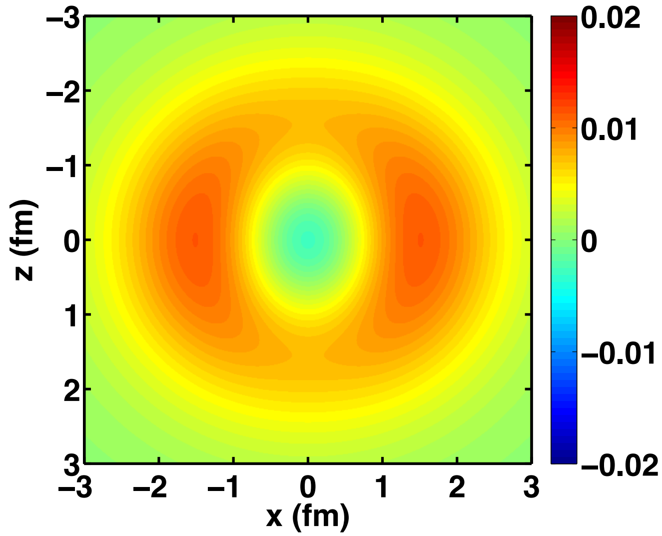

- Relative proton and neutron densities Cockrell, PhD thesis, 2012
- Emergence of α clustering
  - extra neutron appears to be in π orbital

C. Cockrell, J.P. Vary, P. Maris, Phys. Rev. C86, 034325 (2012);
C. Cockrell, PhD, Iowa State University

Predictions for ^{14}F confirmed by experiments at Texas A&M

Theory published PRC: Feb. 4, 2010

Physics Letters B 692 (2010) 307–311

Experiment published: Aug. 3, 2010

ELSEVIER

Contents lists available at ScienceDirect

Physics Letters B

www.elsevier.com/locate/physletb

First observation of ^{14}F

V.Z. Goldberg [a,*], B.T. Roeder [a], G.V. Rogachev [b], G.G. Chubarian [a], E.D. Johnson [b], C. Fu [c], A.A. Alharbi [a,1], M.L. Avila [b], A. Banu [a], M. McCleskey [a], J.P. Mitchell [b], E. Simmons [a], G. Tabacaru [a], L. Trache [a], R.E. Tribble [a]

[a] Cyclotron Institute, Texas A&M University, College Station, TX 77843-3366, USA
[b] Department of Physics, Florida State University, Tallahassee, FL 32306-4350, USA
[c] Indiana University, Bloomington, IN 47408, USA

NCFC predictions (JISP16) in close agreement with experiment

TAMU Cyclotron Institute

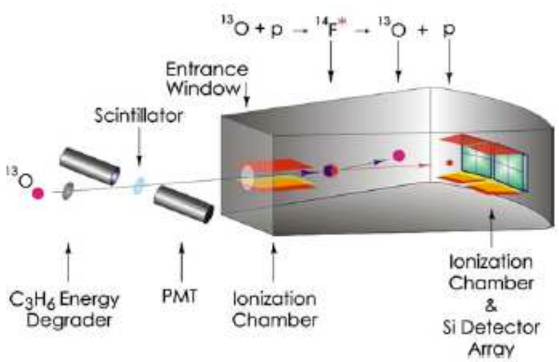

Fig. 1. (Color online.) The setup for the ^{14}F experiment. The "gray box" is the scattering chamber. See explanation in the text.

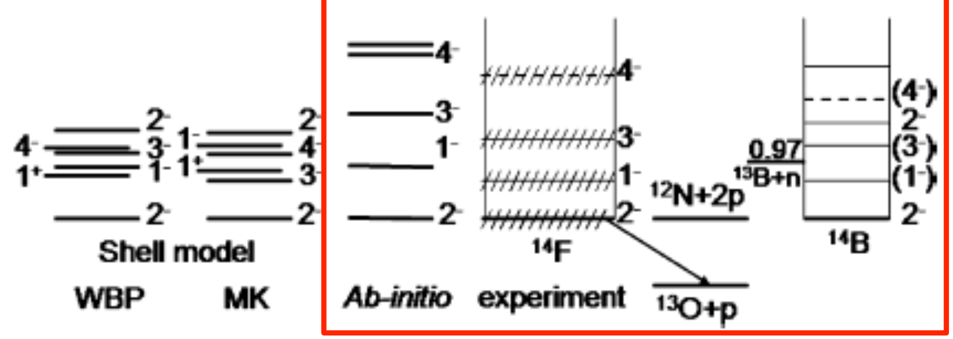

Fig. 6. ^{14}F level scheme from this work compared with shell-model calculations, *ab-initio* calculations [3] and the ^{14}B level scheme [16]. The shell model calculations were performed with the WBP [21] and MK [22] residual interactions using the code COSMO [23].

Ground state magnetic moments with JISP16

Compare theory and experiment for 23 magnetic moments Maris, Vary, IJMPE22, 1330016 (2013)

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

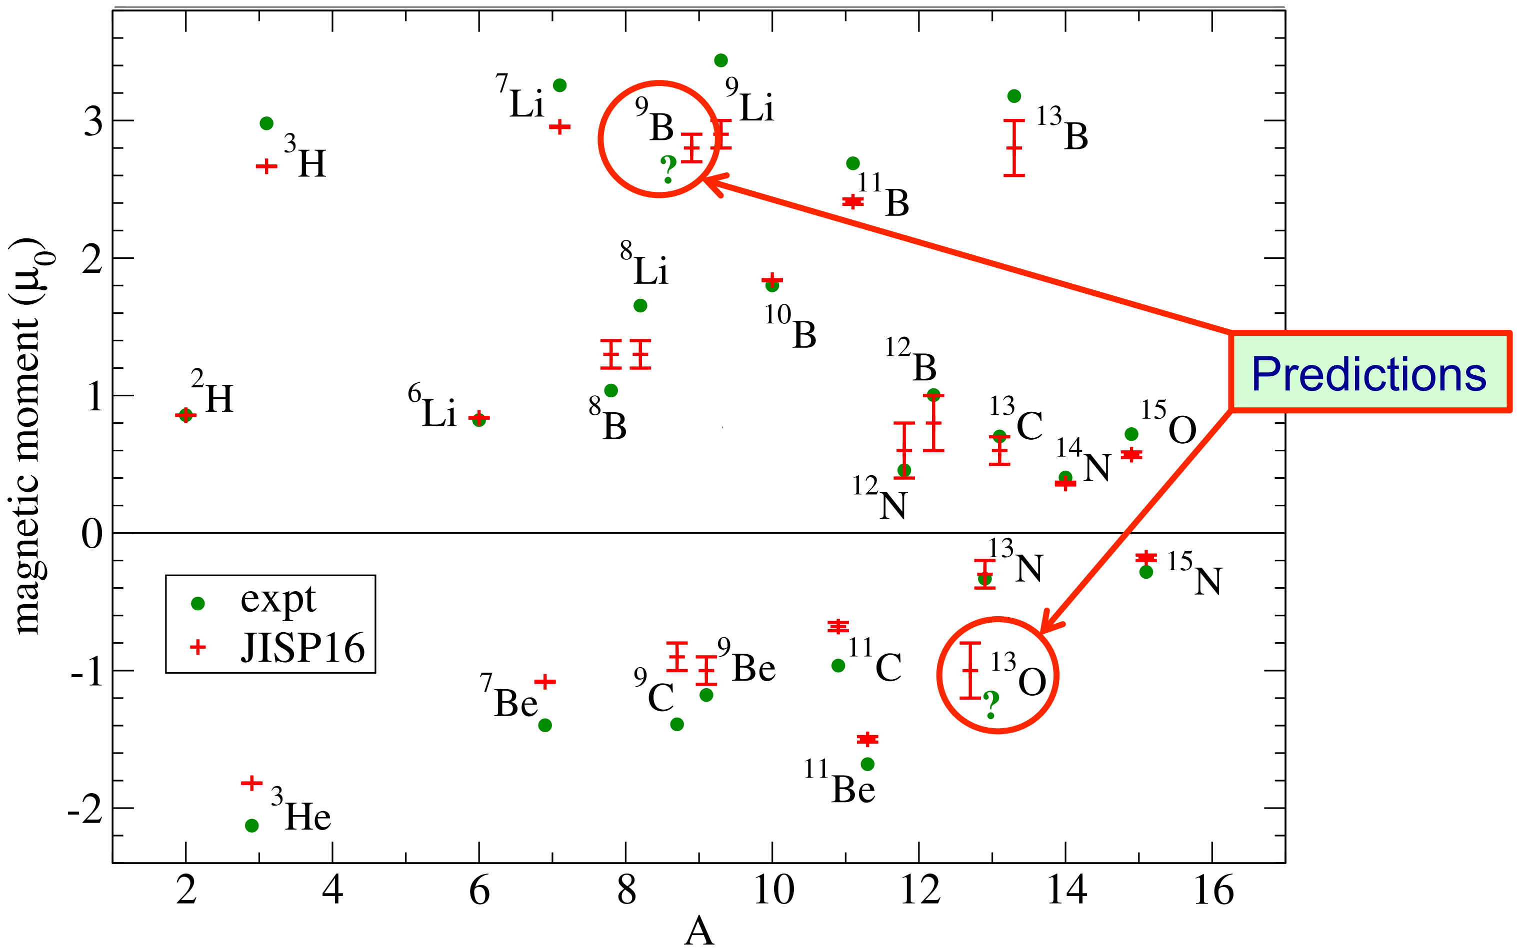

- Good agreement with data, given that we do not have any meson-exchange currents

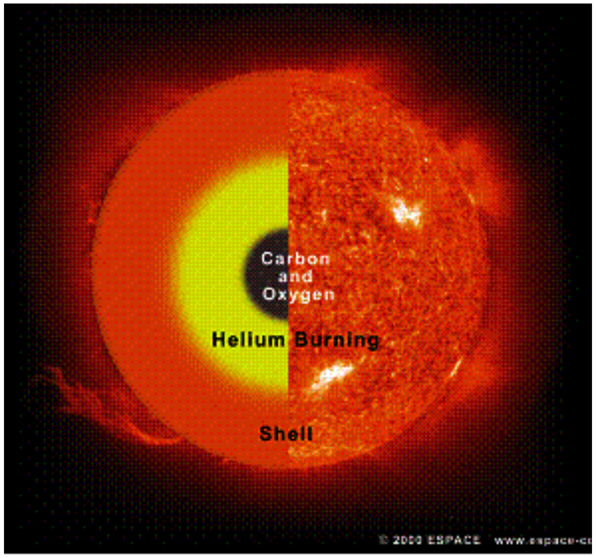

^{12}C - At the heart of matter

The first excited 0+ state of ^{12}C, the "Hoyle state", is the key state of ^{12}C formation in the triple-alpha fusion process that occurs in stars.

Due to its role in astrophysics and the fact that carbon is central to life, some refer to this as one of the "holy grails" of nuclear theory.

Many important unsolved problems of the Hoyle state:

Ab initio origins of the triple-alpha structure are unsolved

Breathing mode puzzle - experiments disagree on sum rule fraction

Laboratory experiments to measure the formation rate are very difficult - resulting uncertainties are too large for predicting the ^{12}C formation rate through this state that dictates the size of the iron core in pre-supernova stars

Conclusion: Need *ab initio* solutions of the Hoyle state with no-core method that accurately predicts the ground state binding energy

==> parameter free predictions for the Hoyle state achievable with petascale within 1-2 years

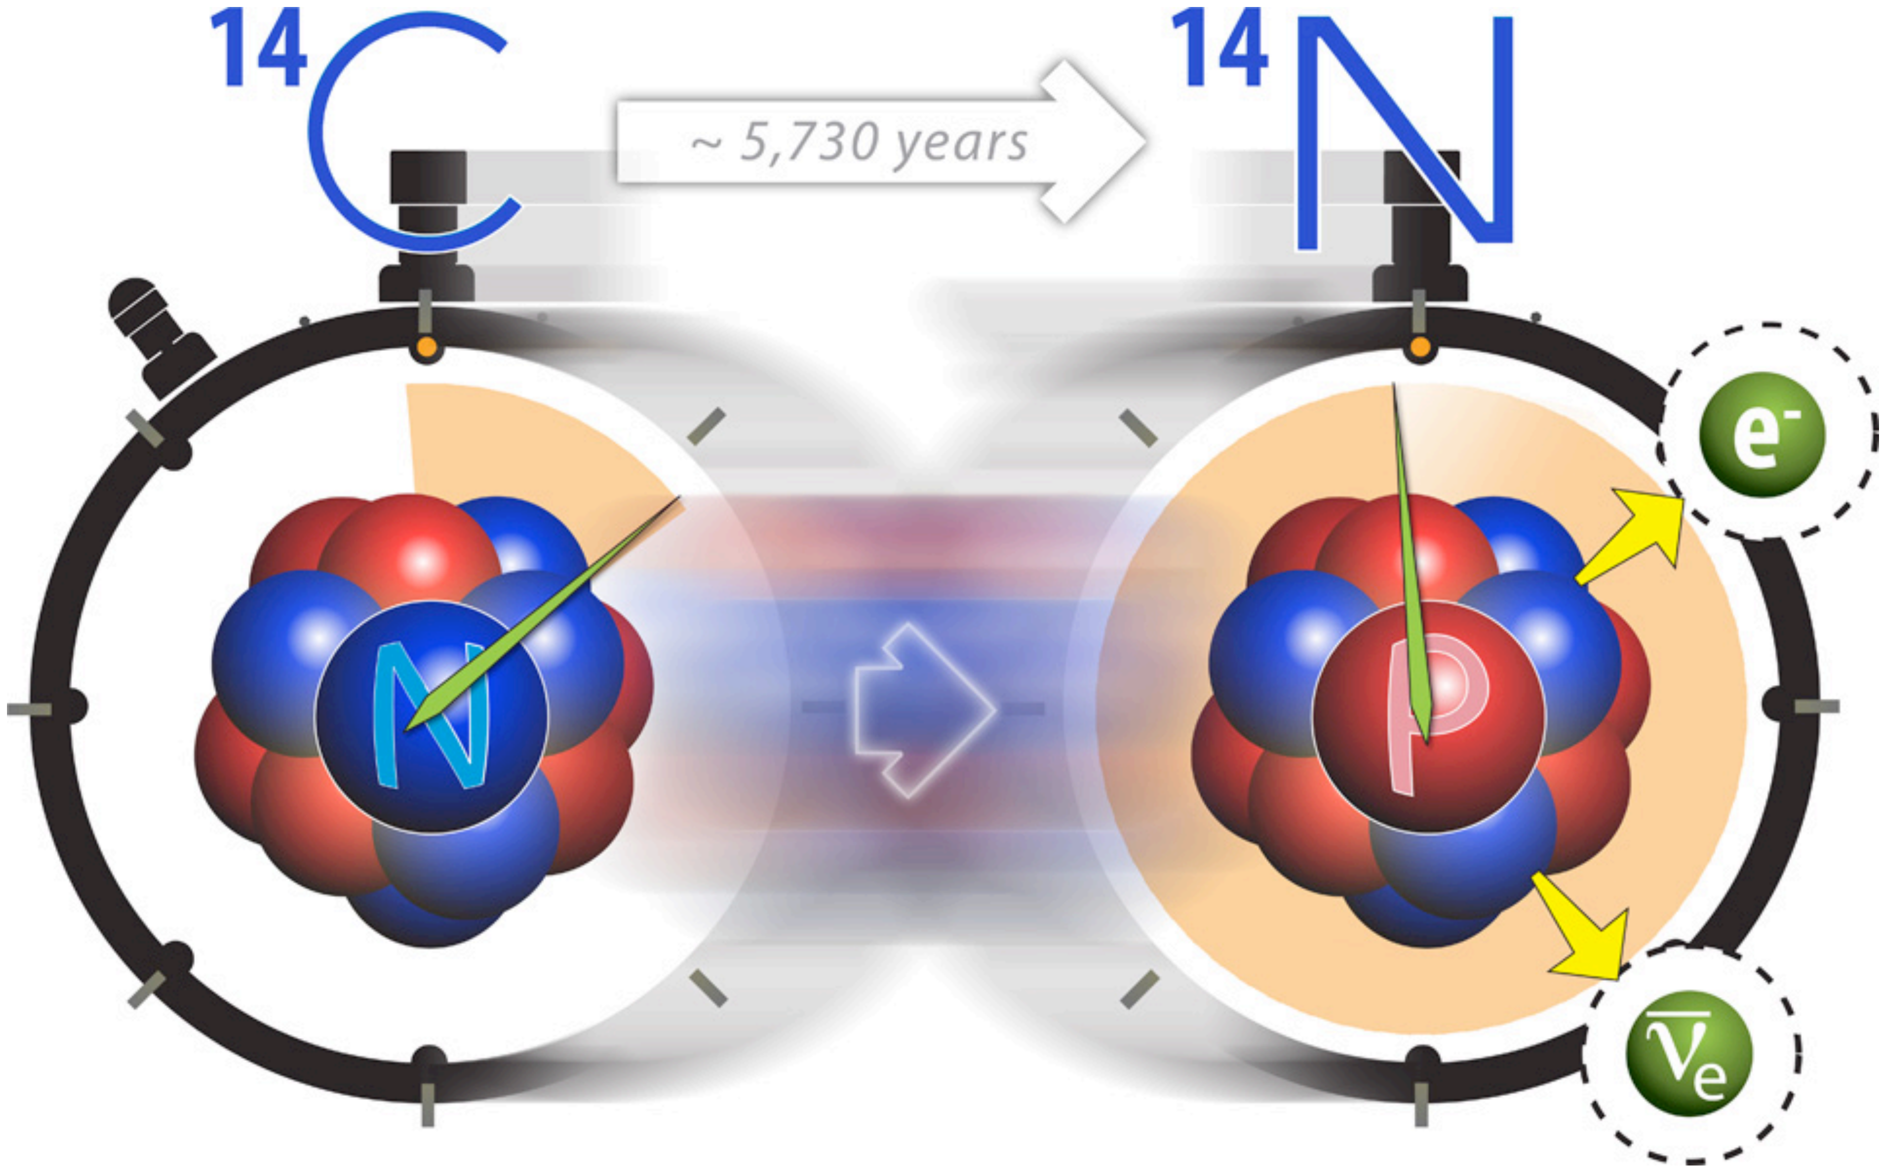

14C
~ 5,730 years
14N
N
p
e-
$\bar{\nu}_e$

"Anomalous Long Lifetime of Carbon-14"

Objectives

- Solve the puzzle of the long but useful lifetime of ^{14}C
- Determine the microscopic origin of the suppressed β-decay rate

Impact

- Establishes a major role for strong 3-nucleon forces in nuclei
- Verifies accuracy of *ab initio* microscopic nuclear theory
- Provides foundation for guiding DOE-supported experiments

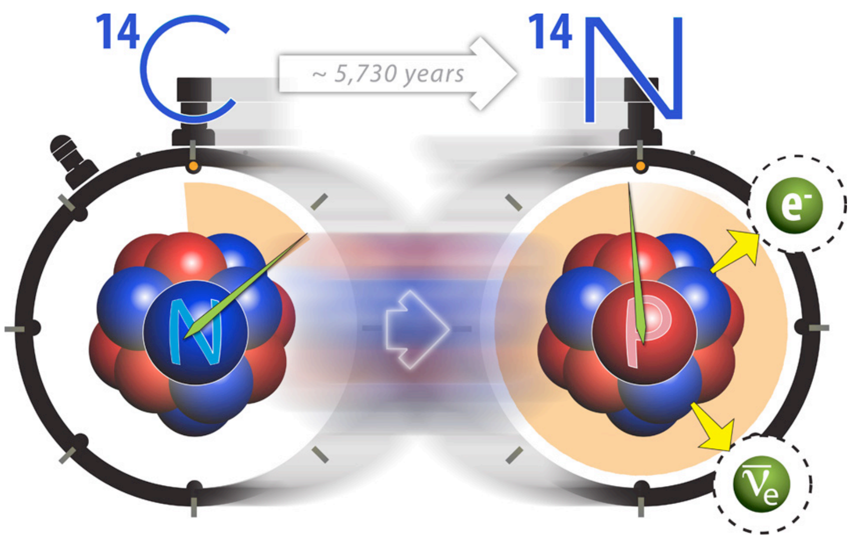

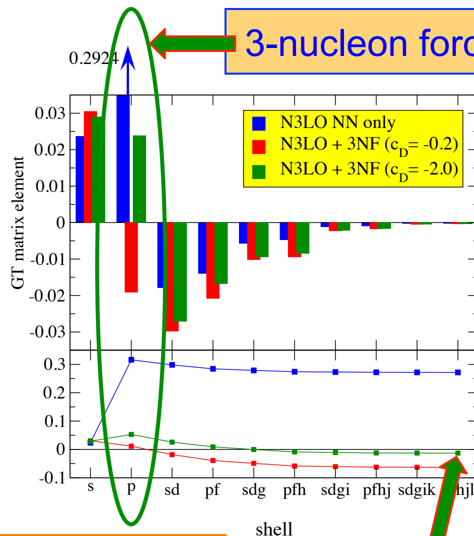

3-nucleon forces suppress critical component

net decay rate
Is very small

- Dimension of matrix solved for 8 lowest states ~ 1×10^9
- Solution takes ~ 6 hours on 215,000 cores on Cray XT5 Jaguar at ORNL
- "Scaling of *ab initio* nuclear physics calculations on multicore computer architectures," P. Maris, M. Sosonkina, J. P. Vary, E. G. Ng and C. Yang, 2010 Intern. Conf. on Computer Science, Procedia Computer Science 1, 97 (2010)

PRL **106**, 202502 (2011) PHYSICAL REVIEW LETTERS week ending 20 MAY 2011

Origin of the Anomalous Long Lifetime of ^{14}C

P. Maris,[1] J. P. Vary,[1] P. Navrátil,[2,3] W. E. Ormand,[3,4] H. Nam,[5] and D. J. Dean[5]

U.S. DEPARTMENT OF ENERGY | Office of Science | INCITE Advancing America's Science and Industrial Competitiveness | IOWA STATE UNIVERSITY | TRIUMF

UNEDF SciDAC Collaboration Universal Nuclear Energy Density Functional | NNSA | LAWRENCE LIVERMORE NATIONAL SECURITY | OAK RIDGE NATIONAL LABORATORY Managed by UT-Battelle for the Department of Energy

Are HPCs in your future?

American Energy and Manufacturing Competitiveness Partnership
2015 Report

High Performance Computing simulations improve product design and speed up development increasing comptitiveness and increasing sales

Figure 2. Complex Layering of Material Used in a Goodyear Tire

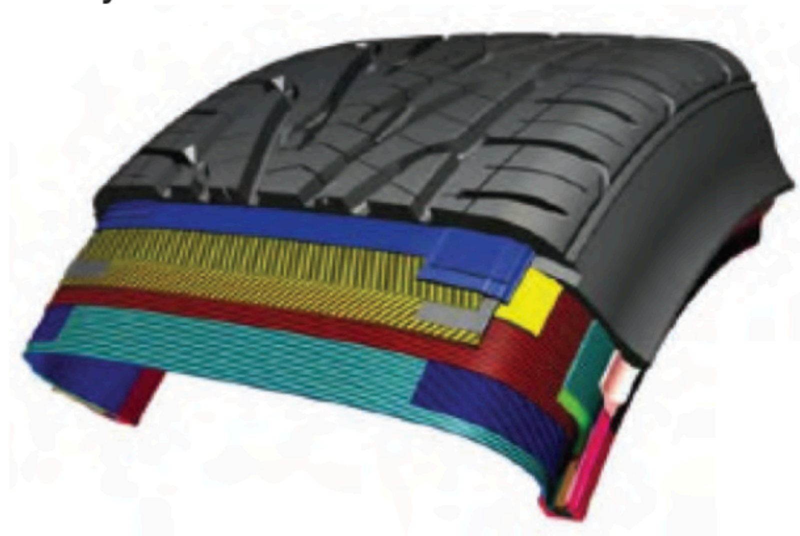

Figure 3. Goodyear Tire & Rubber Company Annual Revenue, 1990-2013

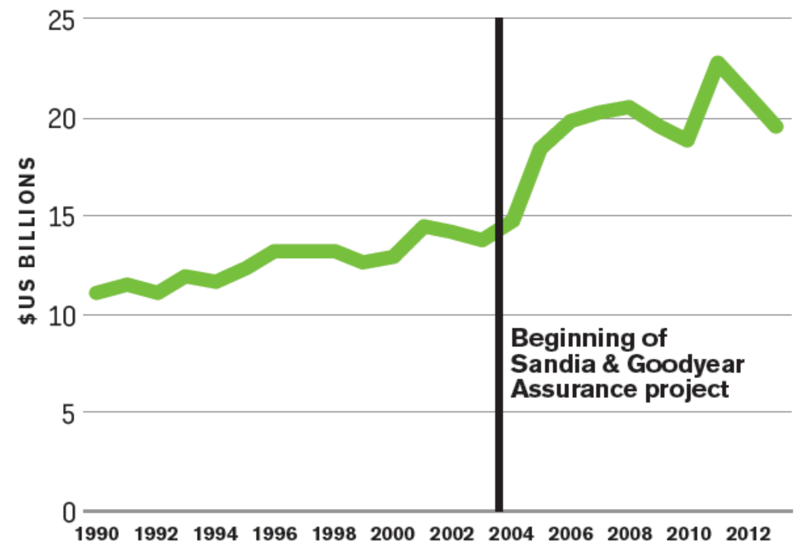

Many additional case studies available:
http://www.compete.org/publications/idea/28/high-performance-computing/

Industrial Partnerships:
Accelerating Competitiveness through Computational Sciences

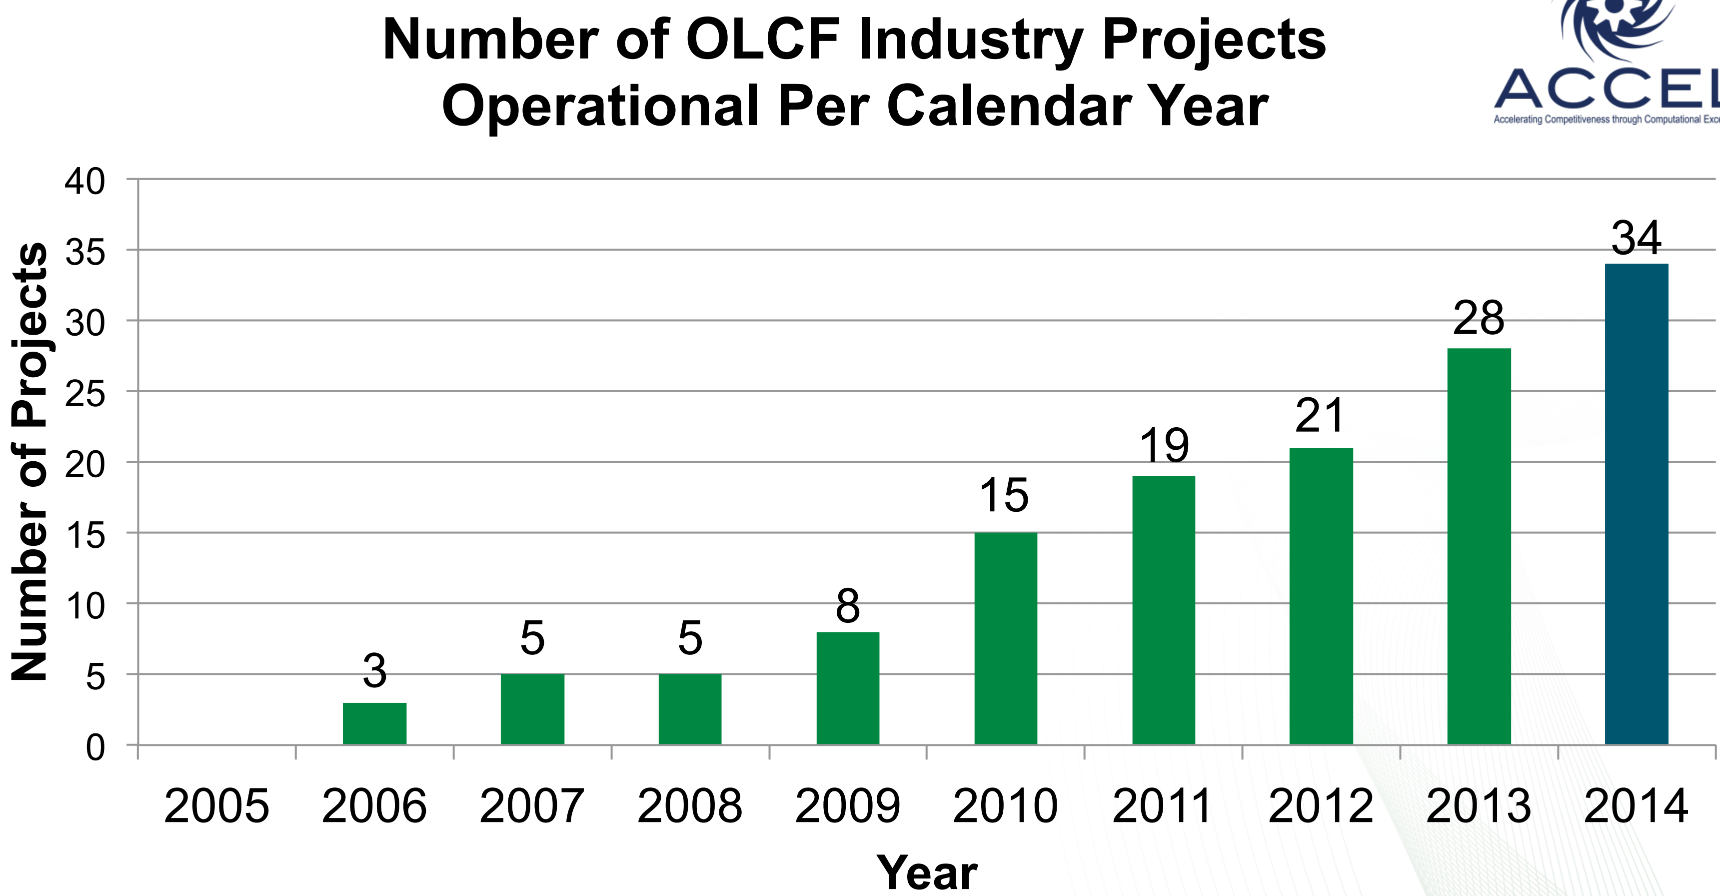

ACCEL
Accelerating Competitiveness through Computational ExceLlence

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Innovation through Industrial Partnerships

BOEING	P&G	GM	United Technologies Research Center	BOSCH	Ford
Aircraft design	**Consumer product stability**	**Gasoline engine injector**	**Jet engine efficiency**	**Li-ion batteries**	**Underhood cooling**
Unexpected discovery of multiple solutions for steady RANS equations with separated flow helps explain why numerical modeling sometimes fails to capture maximum lift	Developed method to measure impact of additives, such as dyes and perfumes, on properties of lipid systems such as fabric enhancer and other formulated products	Optimizing multihole gasoline spray injector nozzle designs for better in-cylinder fuel-air mixture distributions, greater fuel efficiency and reduced physical prototypes	Accurate predictions of atomization of liquid fuel by aerodynamic forces enhance combustion stability, improve efficiency, and reduce emissions	New classes of solid inorganic Li-ion electrolytes could deliver high ionic and low electronic conductivity and good electrochemical stability	Developed a new, efficient and automatic analytical cooling package optimization process leading to one of a kind design optimization of cooling systems
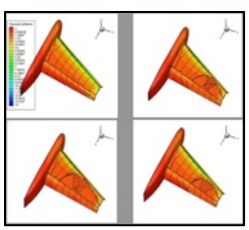	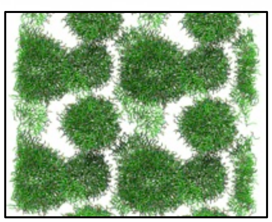	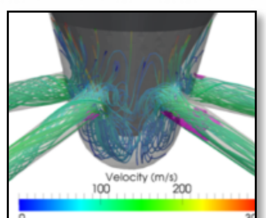	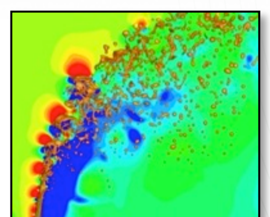	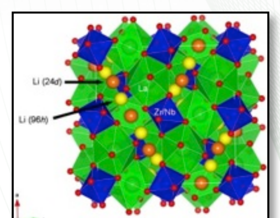	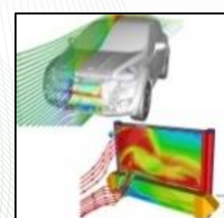

Conclusions and Outlook - I

The tools of science are changing – rise of simulations

“Superfacilities for Science”

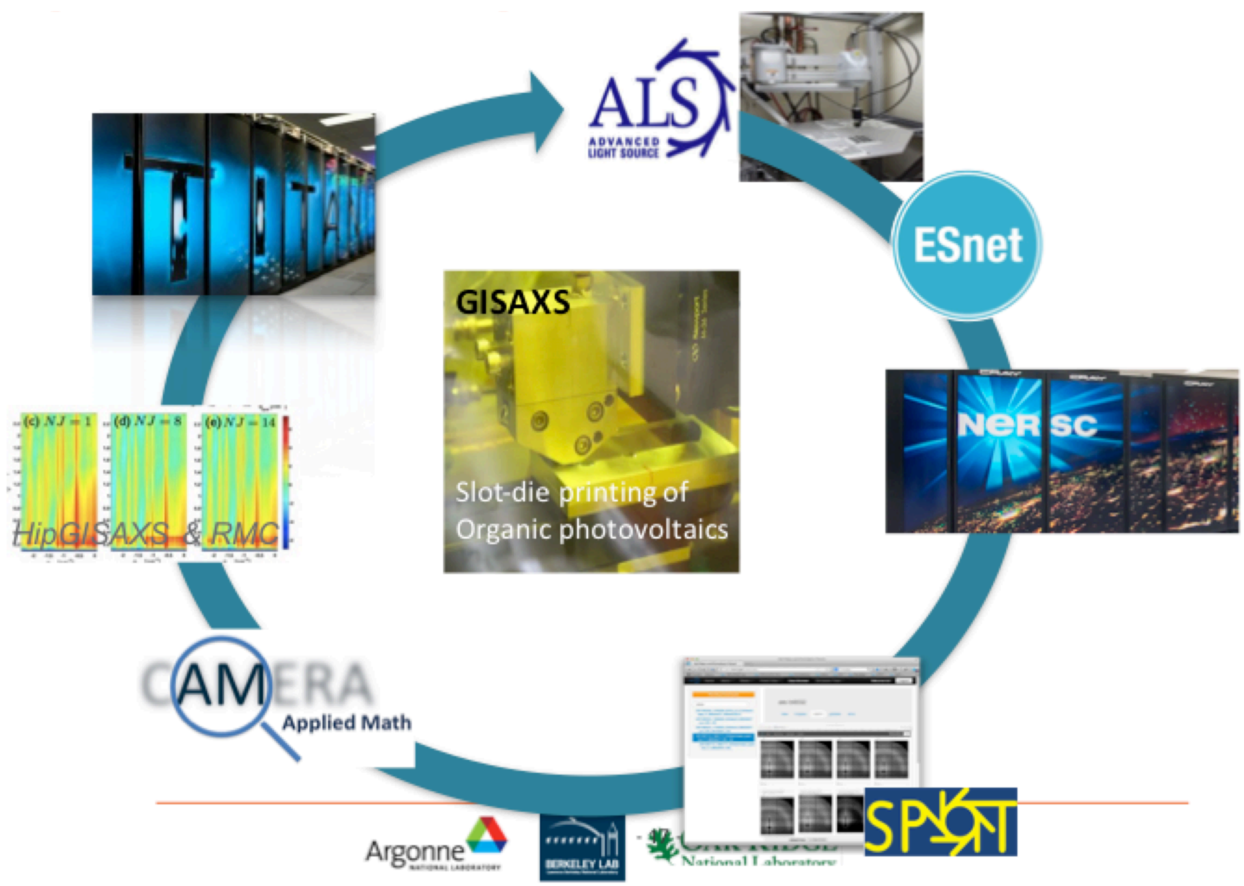

Conclusions and Outlook - II

Many challenges exist to use these tools effectively

MFDn's comm. graph

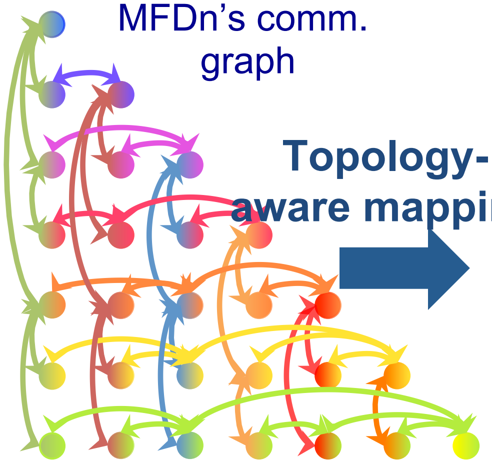

Topology-aware mapping

3D torus

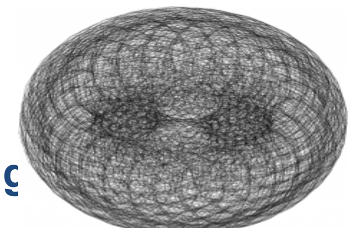

Communication Hiding

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 [1,2].

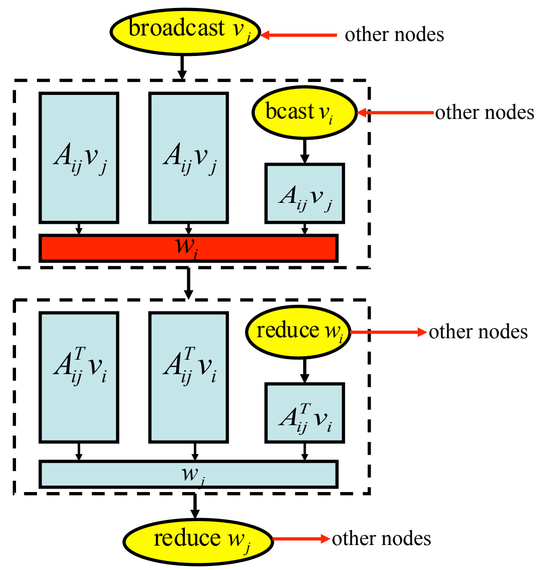

H.M. Aktulga, C. Yang, E.G. Ng, P. Maris, J.P. Vary, "Improving the Scalability of a Symmetric Iterative Eigensolver for Multi-core Platforms", Concurrency Computat.: Pract. Exper., 25 (2013)
D. Oryspayev, H.M. Aktulga, M. Sosonkina, P. Maris and J.P. Vary, "Performance Analysis of Distributed Symmetric Sparse Matrix Vector Multiplication Algorithm for Multi-core Architectures," Concurrency Computat.: Pract. Exper., 27, 5019 (2015)

Conclusions and Outlook - III

Promise of a new era of science – Discoveries with Simulations

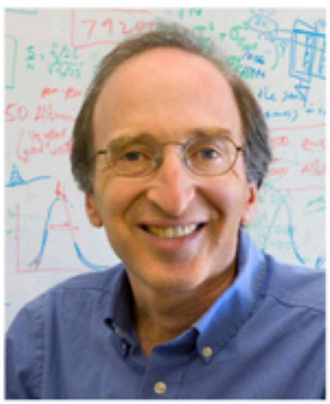

Saul Perlmutter—a professor of physics at UC Berkeley and a faculty senior scientist at Berkeley Lab—was awarded the 2011 Nobel Prize in Physics for his 1998 discovery that the universe is expanding at an accelerating rate. He confirmed his observations by running thousands of simulations at NERSC, and his research team is believed to have been the first to use supercomputers to analyze and validate observational data in cosmology.