Computational nuclear physics…

Input for an LRP initiative

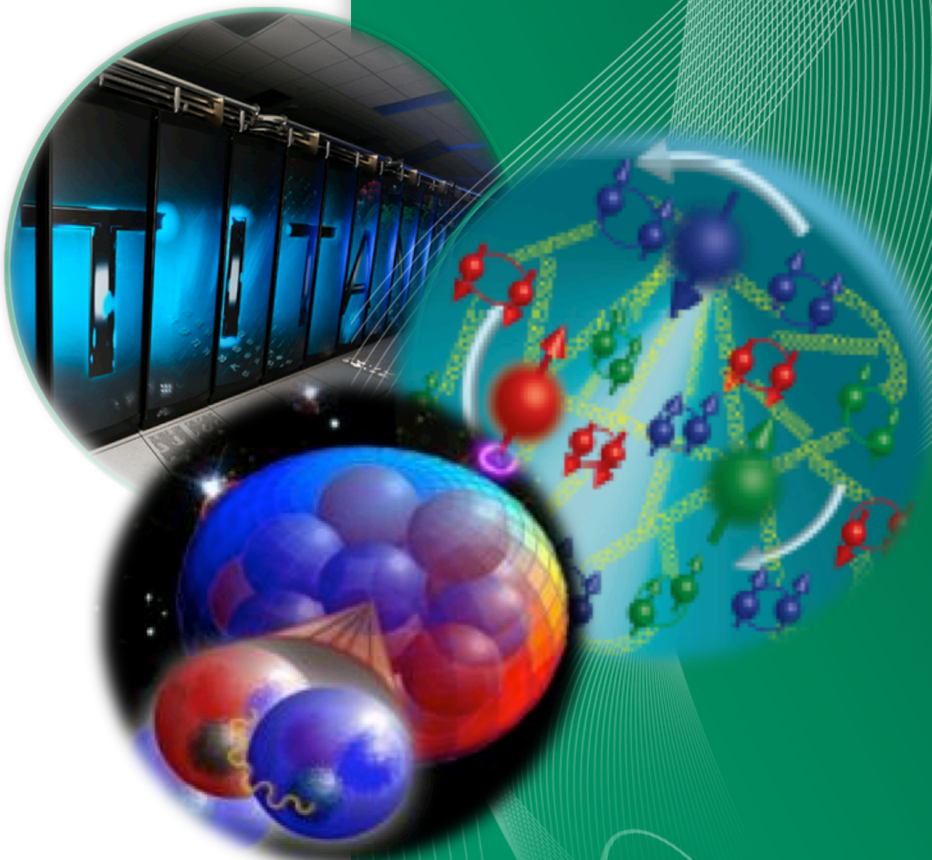

David J. Dean

Director, Physics Division

ORNL

HPC meeting

Washington, DC

July XX, 2014

Outline

- Quick scientific overview
  - Representative review of where we thought we would be in 2009 (via the ASCR/NP report) compared to where we are today (2014).
- Outline the direction of ASCR computing with a quick look at the evolution of the top500, and the likely trajectory during the next 5-10 years.
- Progress of our colleagues:
  - Recent HEP report on computing and what they think is important and why
  - Current direction of BES and OSTP in the 'materials genome initiative'.
- What should we propose to LRP (discussion)?

From small-x to heavy nuclei…

Scientific principles and laws do not lie on the surface of nature. They are hidden, and must be wrested from nature by an active and elaborate technique of inquiry. ~John Dewey, *Reconstruction in Philosophy*, 1920

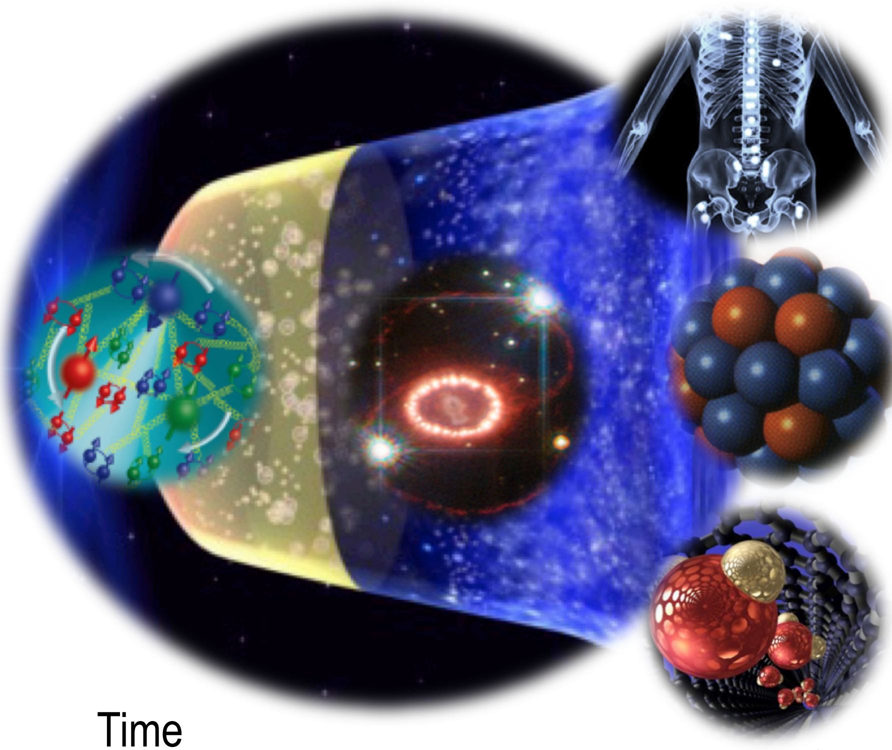

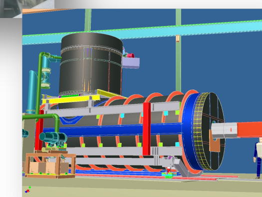

Applications of Nuclei

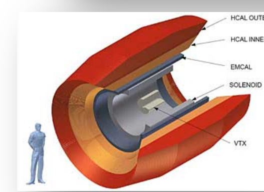

Unity of Science

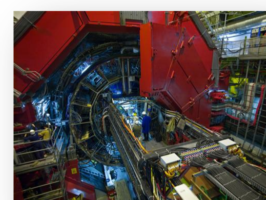

Time

Science is always wrong. It never solves a problem without creating ten more. ~George Bernard Shaw

Intellectual links in field…

QCD

From S. Aoki

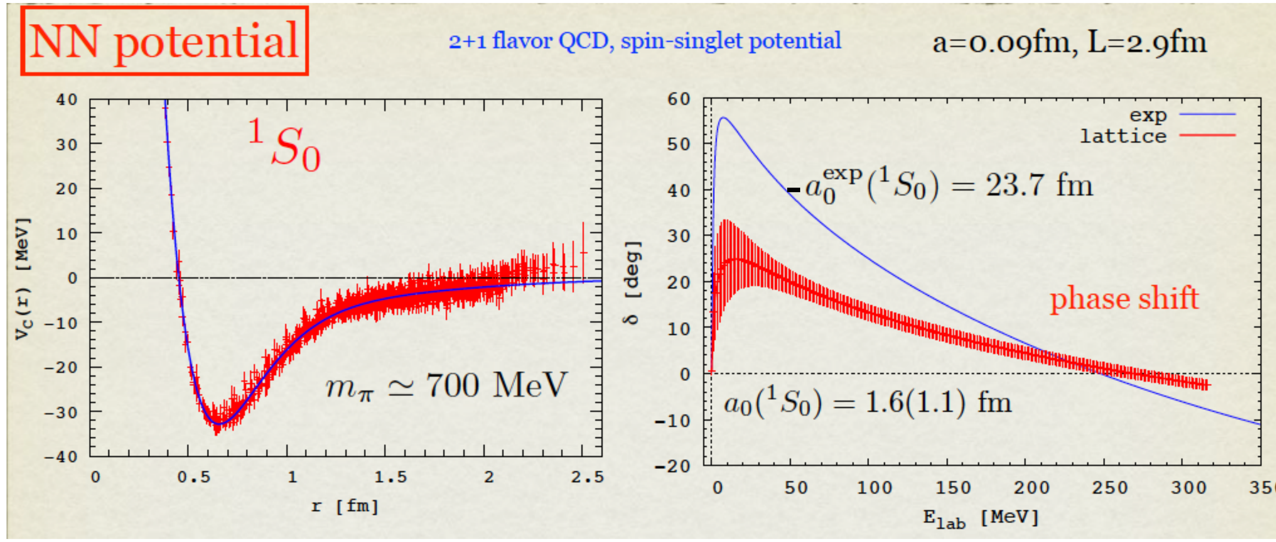

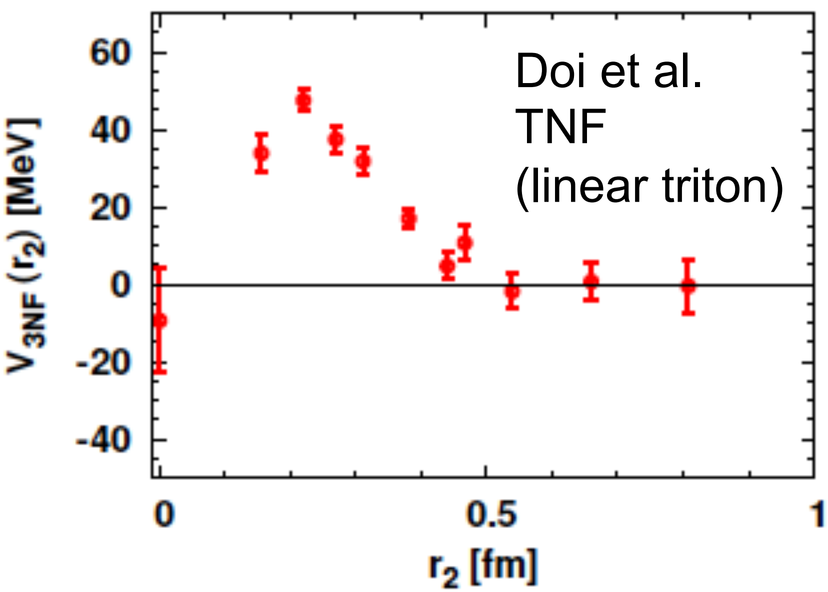

Lattice QCD and the NN (and 3N) interaction may one day be fully linked…

Until then…
Chiral Effective Field Theory…

Progress in calculating nuclei

Moving toward predictive capability

EFT + Continuum + HPC → Neutron rich nuclei from theory

+

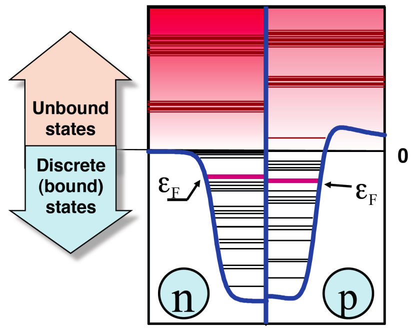

+

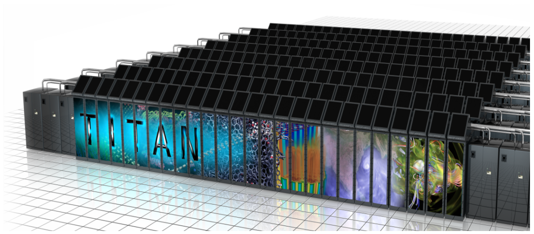

=

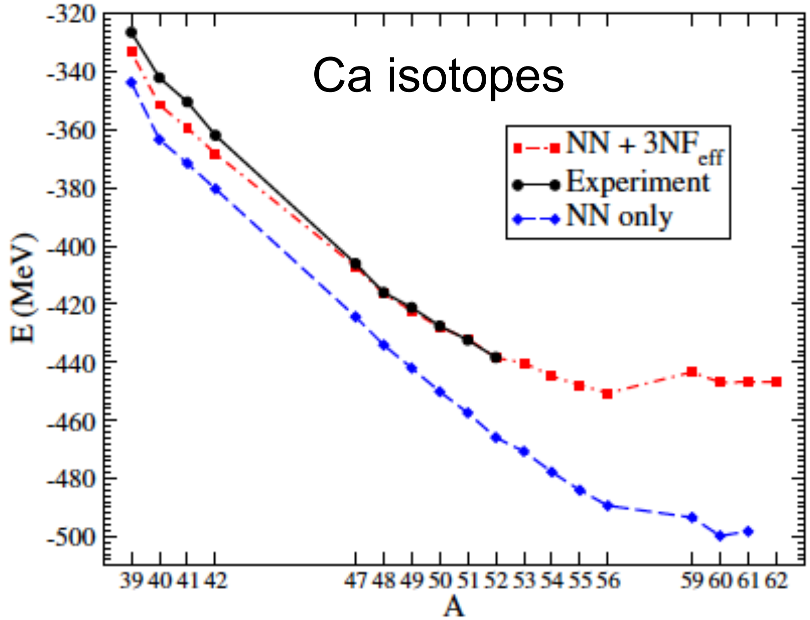

Hagen et al. PRL 109, 032502 (2012)

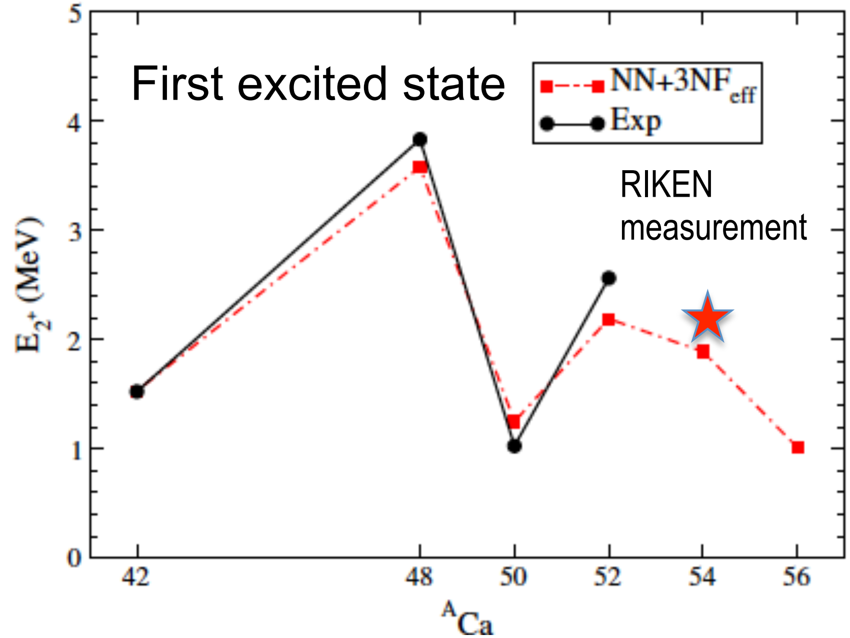

Current status:
Improving EFT interactions in nuclei and matter…

Recent CC results: Optimized NN

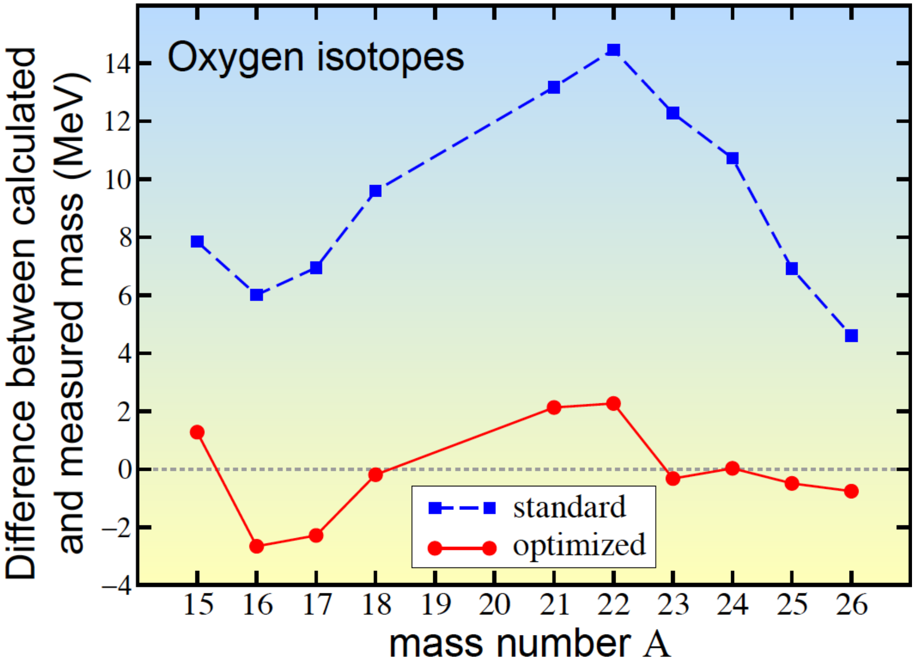

- POUNDERS optimization of NN chiral force (replaces χ-by-eye)
- developed in UNEDF SciDAC project
- Coupled-cluster theory results
- Ekstrom et al, PRL 110, 192502 (2013)

BUT…

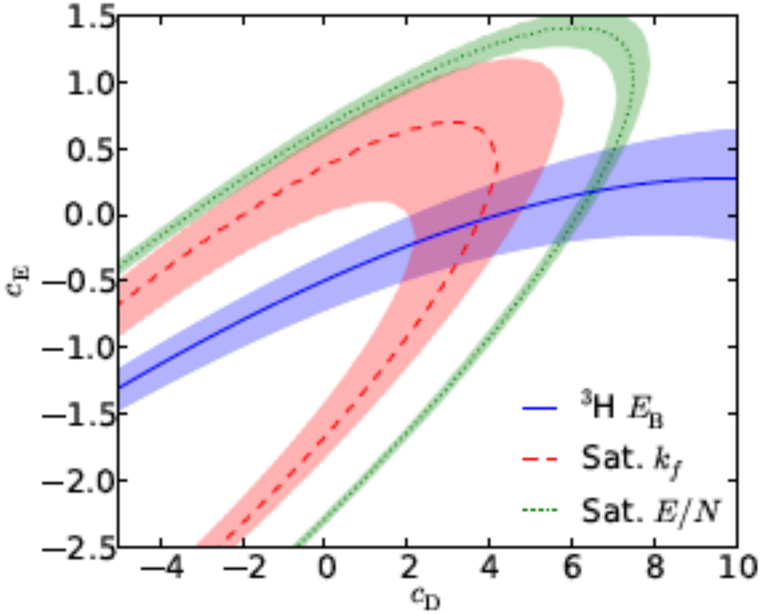

- Excellent neutron matter saturation results
- Symmetric nuclear matter still presents a challenge: Hagen et al (arXiv:1311.2925)

Current Status:
Nuclear Density Functional Approach

- Full quantum many-body approaches (GFMC, NCSM, CC) based on NN+3N cannot tackle every nucleus
- Density Functional Theory: if you know the energy density functional, you can precisely determine the ground-state properties of a quantum many-body system (Kohn-Sham theorem – Nobel, 1998)

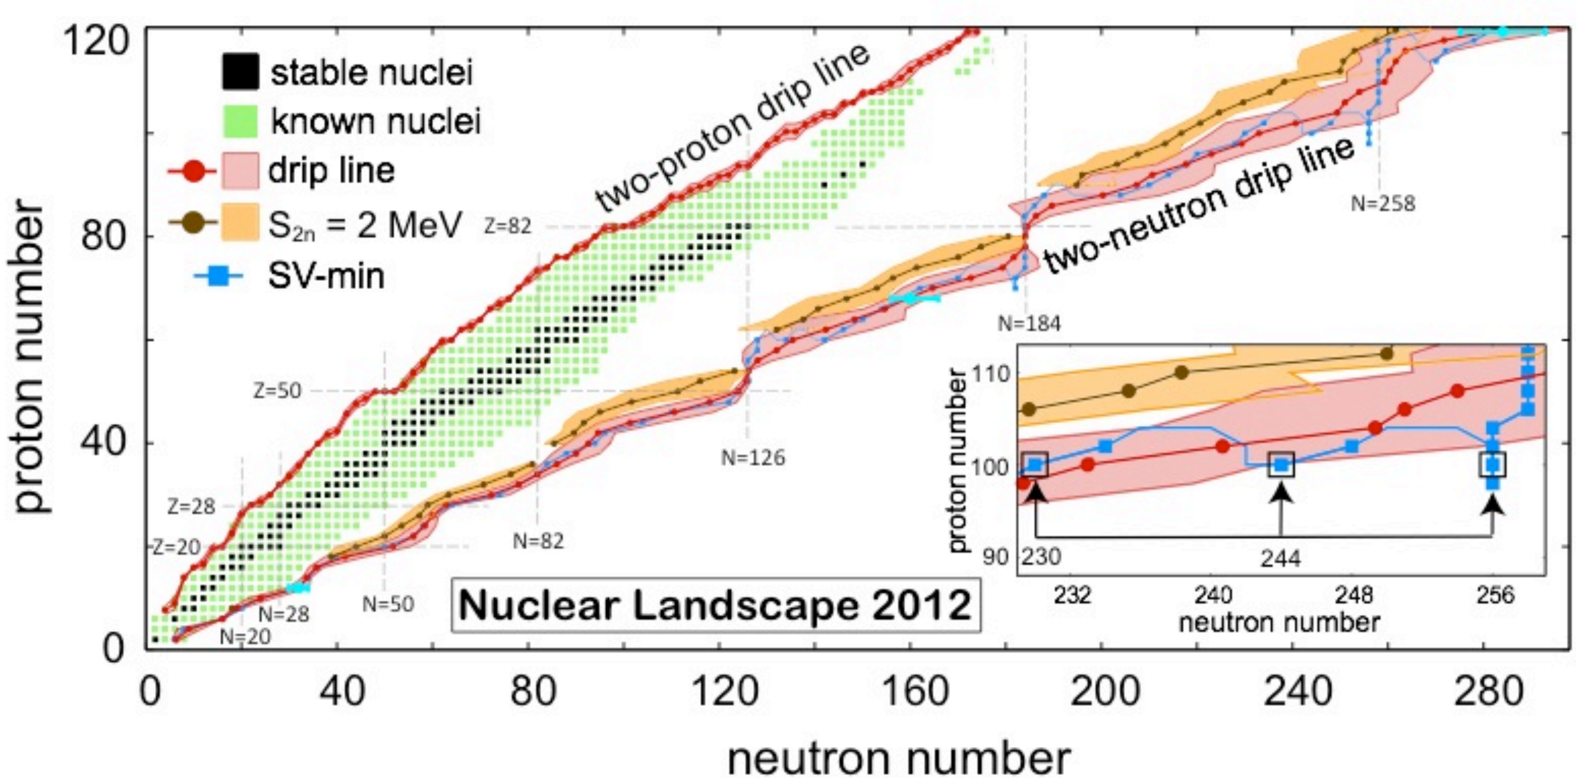

Number of bound nuclei

6900 ± 500

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

Astrophysics connections

How did visible matter come into being and how did it evolve?

X-ray burst

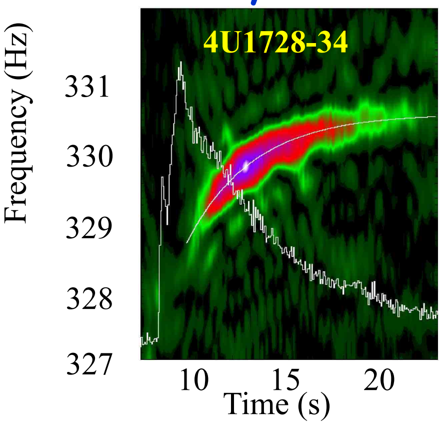

Supernova

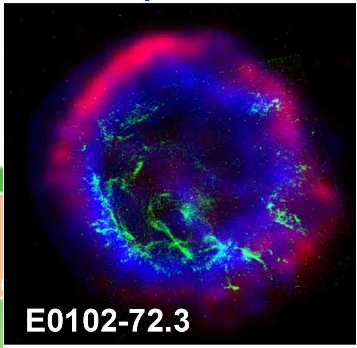

Nova

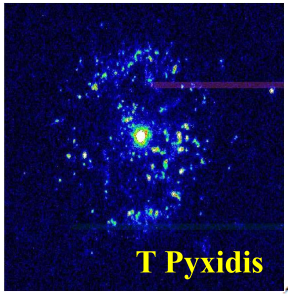

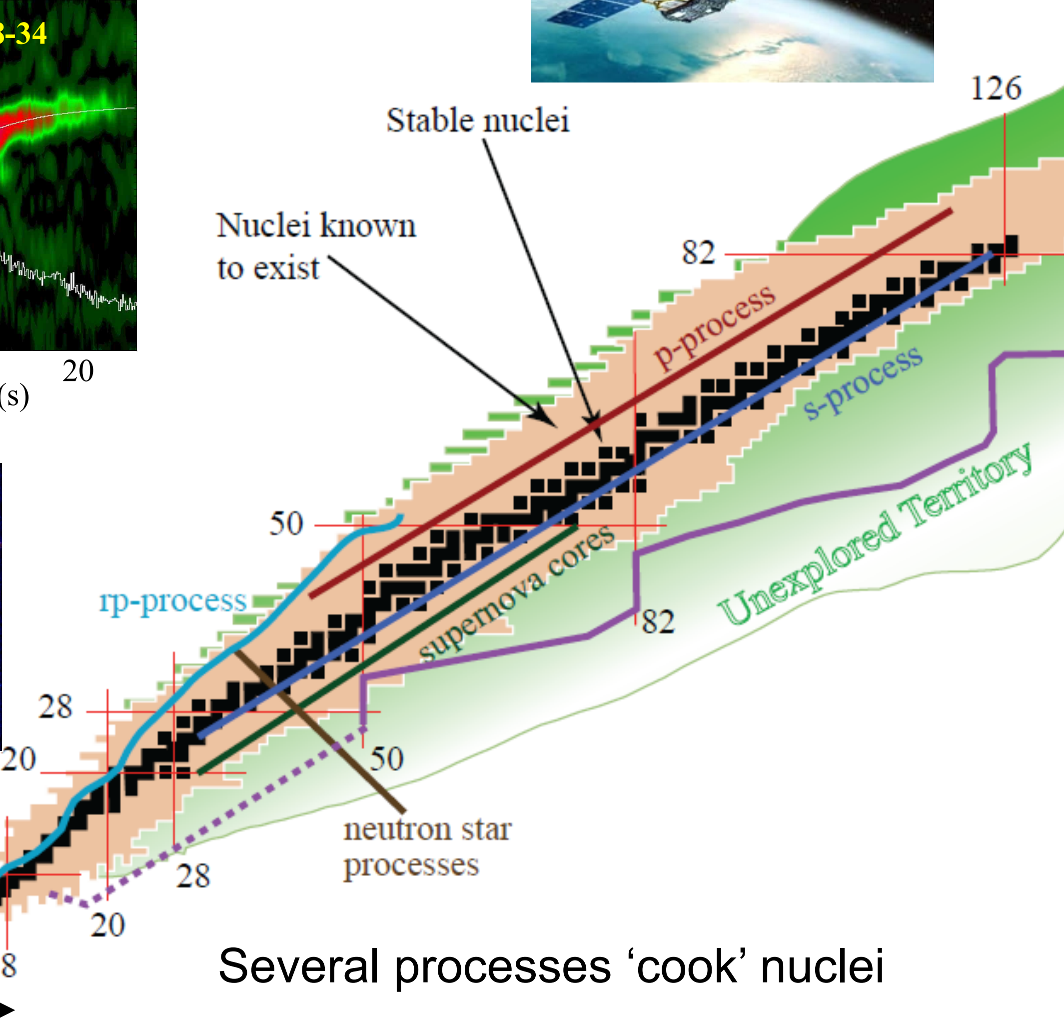

Neutron star

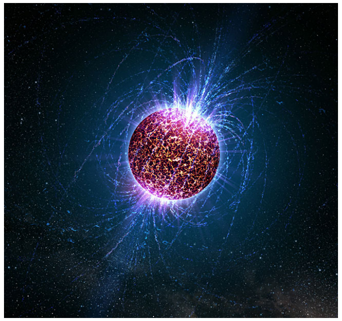

Several processes ‘cook’ nuclei

R-process and SN (or mergers)

How did visible matter come into being and how did it evolve?

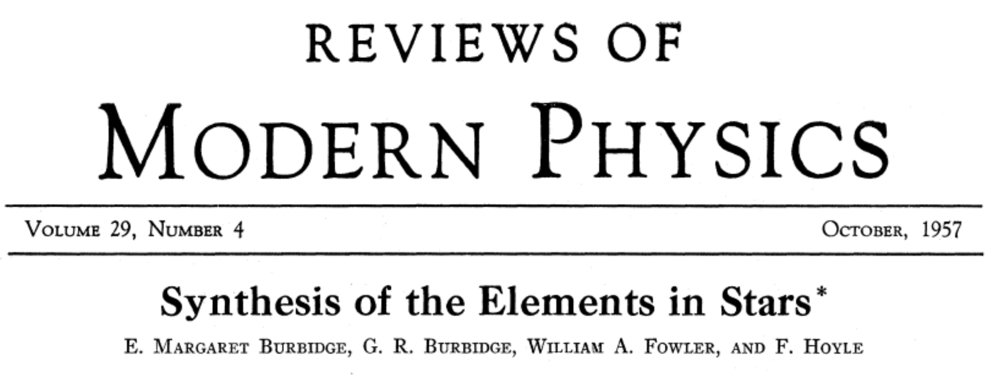
REVIEWS OF

MODERN PHYSICS

VOLUME 29, NUMBER 4 — OCTOBER, 1957

Synthesis of the Elements in Stars*

E. MARGARET BURBIDGE, G. R. BURBIDGE, WILLIAM A. FOWLER, AND F. HOYLE

One example:

R-process: rapid neutron capture responsible for ½ of the heavy elements

Requires:

- Neutron density: $10^{20\text{-}28}$ n/cm^3
- Fast time scale (seconds)
- Astrophysical site unknown

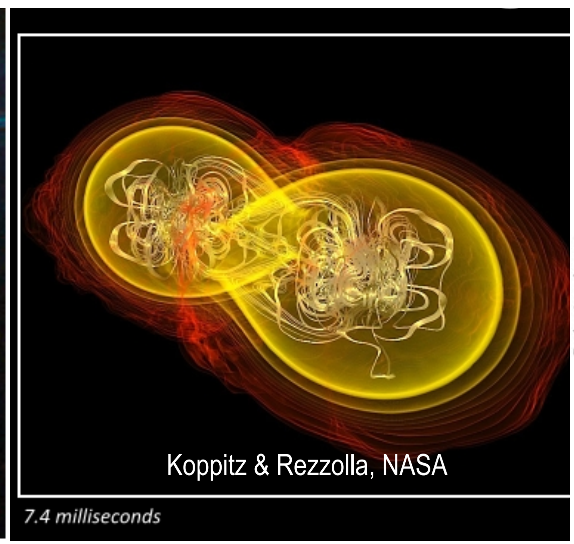

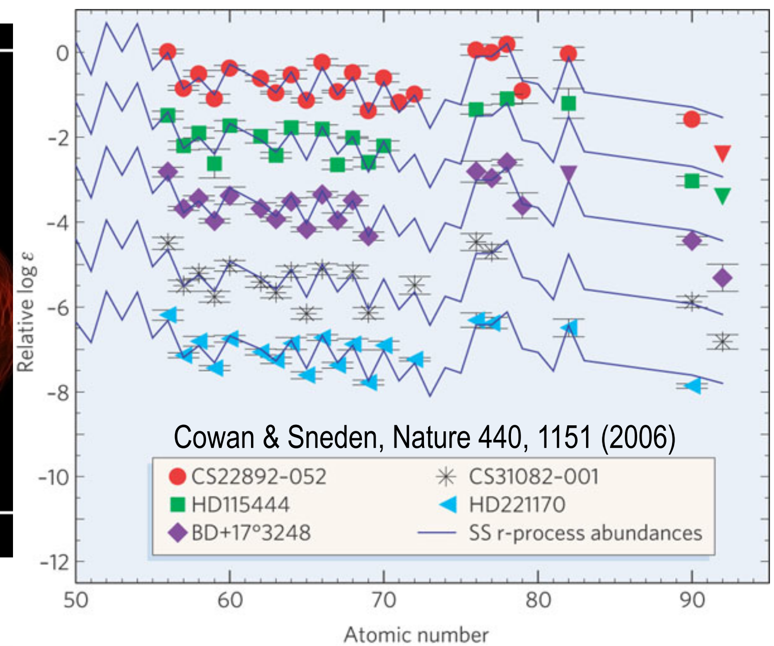

HEP and NP OLCF projects (2014)

Alloc Type	ProjectID	PI	PI Employer	Project Name	Alloc
ALCC_2014	NPH014	Aurel Bulgac	University of Washington	Structure and Dynamics of Nuclear Systems within Time-Dependent Density Functional Theory Approach	25,000,000
ALCC_2015	NPH102	Keh-Fei Liu	University of Kentucky	Quark and Glue Structure of the Nucleon with Lattice QCD	68,800,000
ALCC_2015	NPH103	Robert Glenn Edwards	JLab	The Spectrum and Properties of Exotic Mesons in Quantum Chromodynamics	250,000,000
ALCC_2015	NPH104	Martin Savage	University of Washington	Hypernuclei and Charmed Nuclei	65,100,000
DD_2014	NPH013	Dipangkar Dutta	Mississippi State University	A New Search for the Neutron Electric Dipole Moment	1,470,000
DD_2014	NPH015	Kenneth Read	ORNL	Probing Fluctuating Intitial Conditions of Heavy-Ion Collisions	300,000
DD_2014	NPH101	Jirina Rikovska Stone	University of Oxford	Phase transitions in high density matter in neutron stars and supernovae	8,000,000
DD_2014	CSC108	Sergey Panitkin	BNL	Next Generation Workload Management System	10,500,000
INCITE_2014	NPH008	James Vary	ORNL	Nuclear Structure and Nuclear Reactions	104,000,000
INCITE_2014	LGT003	Paul Mackenzie	FNAL	Lattice QCD	100,000,000
					633,170,000
DD_2014	AST014	Bronson Messer	ORNL	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae	6,000,000
DD_2014	AST032	Simon Portegies Zwart	Leiden University	The Fine Structure of the Milky Way Galaxy	3,000,000
DD_2014	AST103	Michael Andrew Clark	NVidia	Petascale Cross Correlation	2,000,000
DD_2014	AST104	Alexei Kritsuk	University of California	High-resolution Simulations of Compressible MHD turbulence on GPU	3,000,000
DD_2014	AST105	Dominique Aubert	University Strasbourg	BEMMA : Benchmarking Emma	2,000,000
INCITE_2013	AST031	Paul Shapiro	University of Texas Austin	Simulating Reionization of the Local Universe: Witnessing our own Cosmic Dawn	40,000,000
INCITE_2014	AST005	Eric Lentz	ORNL	Three-dimensional simulations of core-collapse supernovae with Chimera	85,000,000
INCITE_2014	AST006	Stan Woosley	University of California Santa Cruz	Petascale Simulations of Type Ia Supernovae	50,000,000
INCITE_2014	AST102	Michael Warren	LANL	Probing Dark Matter at Extreme Scales	80,000,000
					271,000,000

HE+NP = 904M core hours = 1,605 wall clock hours across machine = 18% of machine = $16M/year leverage

Have we accomplished 0.1 – 1.0 PF year milestones from 2009 ASCR/NP report?

- Roughly speaking: the community has about 0.5 PF-year sustained effort on Titan (probably equivalent on ALCF and NERSC; assumes 100% efficiency…

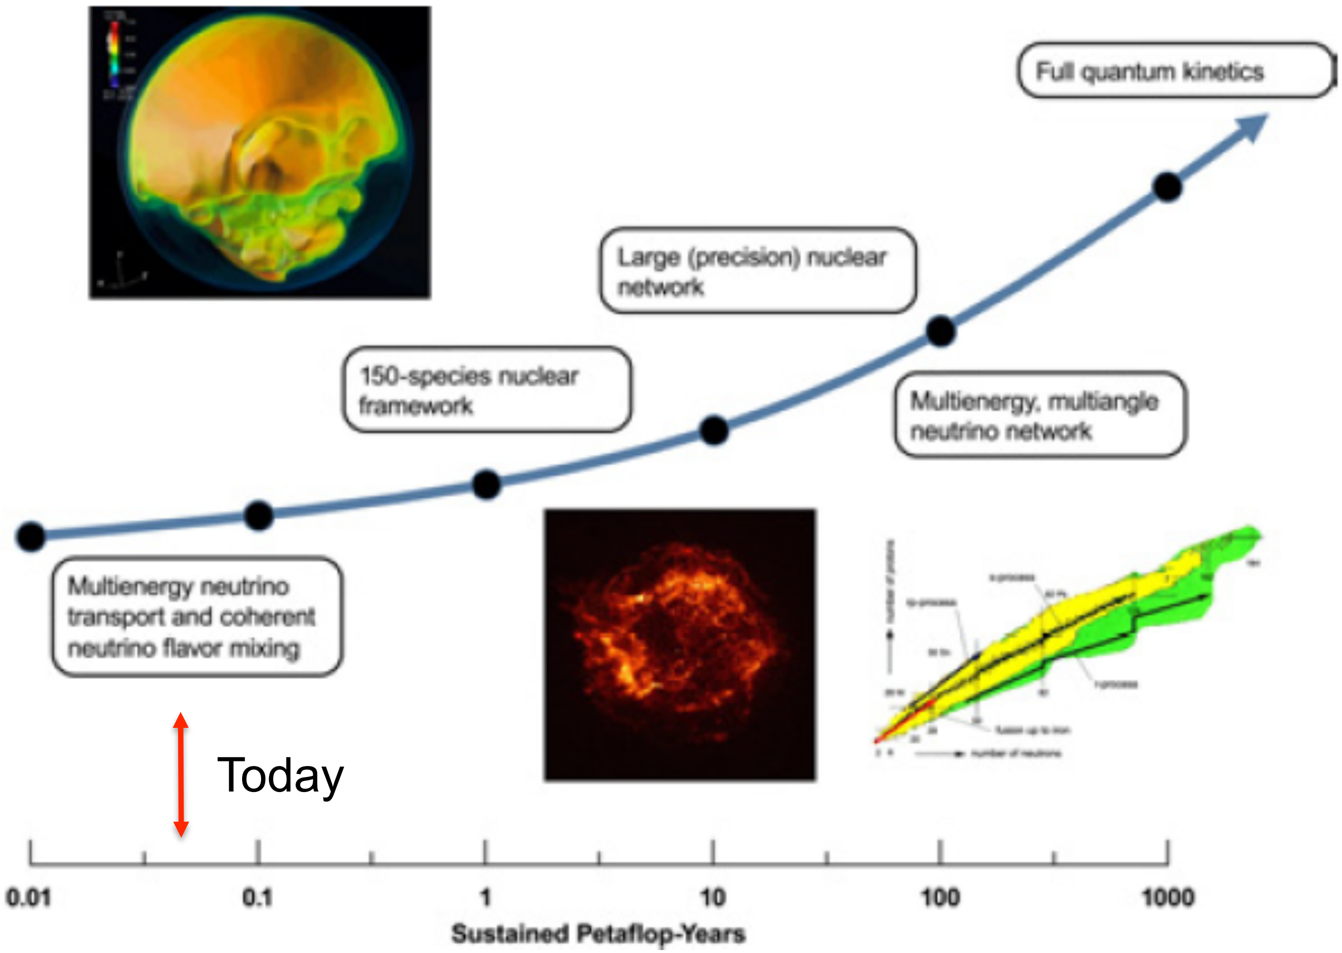

Physics of Nuclei assessment…

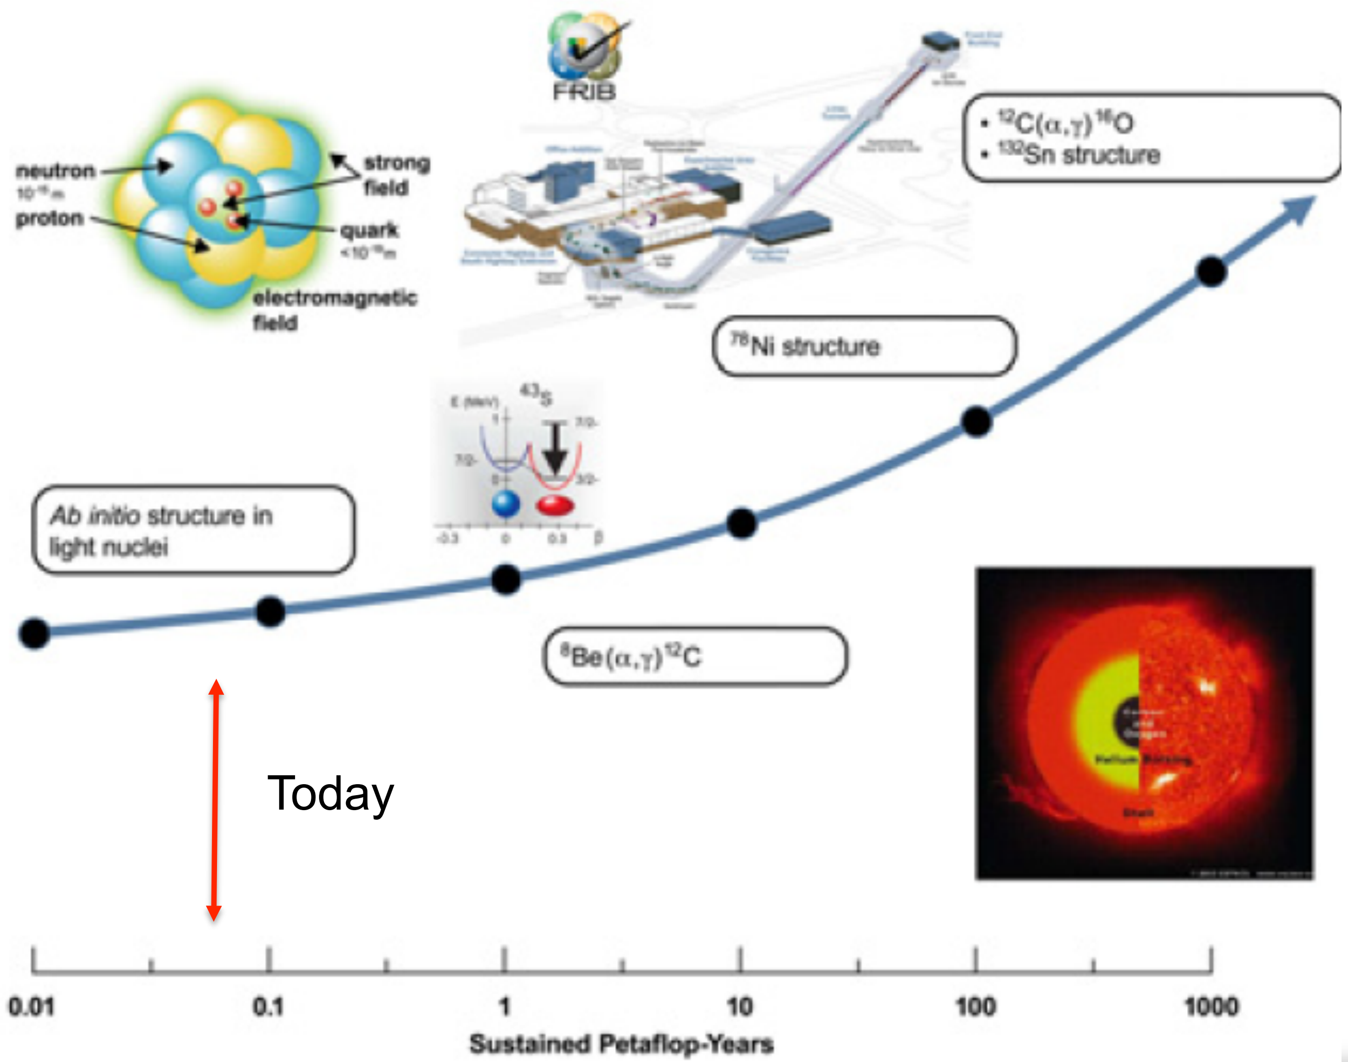

Top 500 (June, 2014)

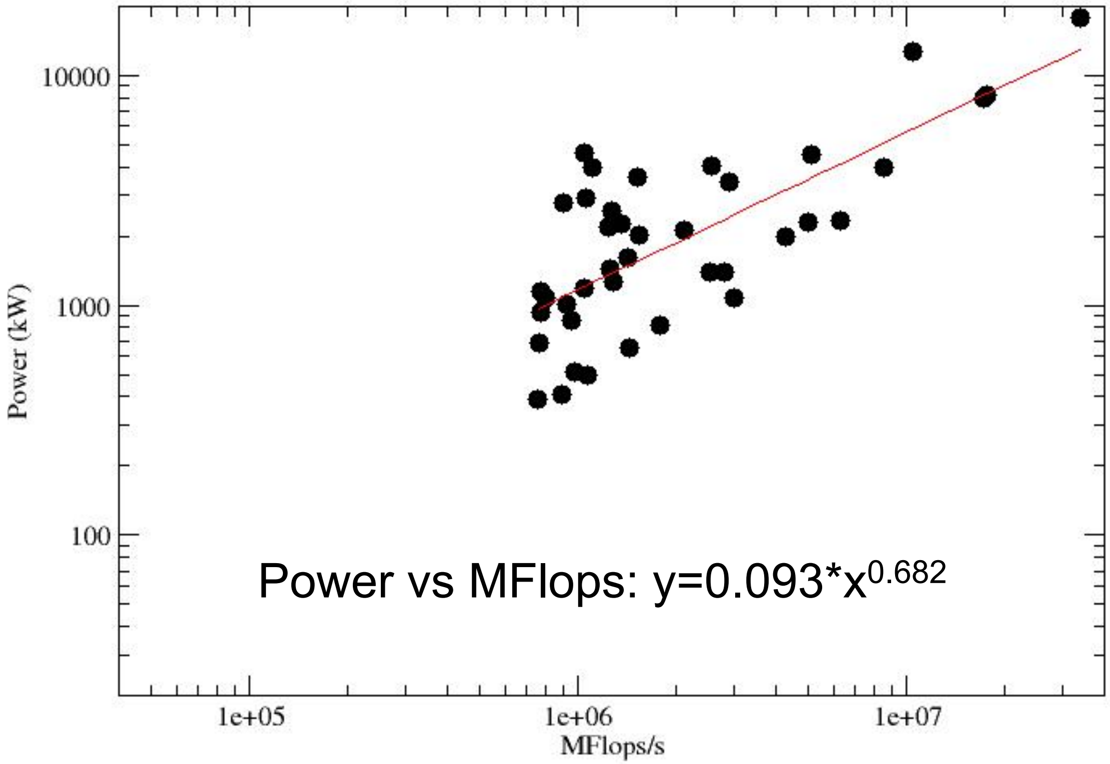

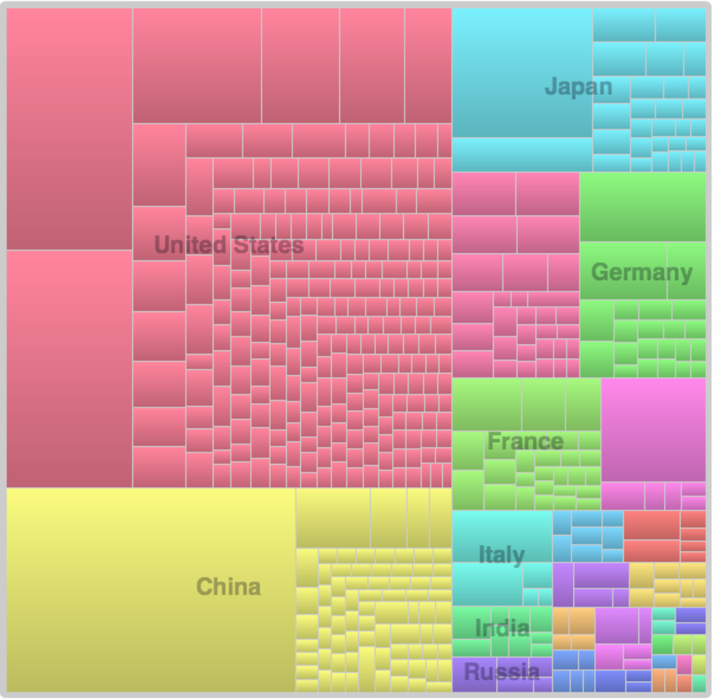

Expertise buildup worldwide

Projections in 2009: 1 Ef machine = 1 GW

TODAY
Power law fit: 300 PF machine = 60 MW power; 1 EF machine = 131 MW power
Drastic improvements in power usage (and accelerators)

1 PF machine (still substantial architecture and infrastructure!):
$2-5M/year infrastructure+replacement

The OLCF 10-year plan

- OLCF has a 10-year plan to deploy and operate the computational resources required to tackle science problems of global importance

	2012	2016	2020	2024
Peak flops	10-20 PF	100-200 PF	500-2000 PF	2000 - 4000 PF
Memory	0.5-1 PB	5-10 PB	32 – 64 PB	50-100 PB
Burst storage bandwidth	NA	5 TB/s	32 TB/s	50 TB/s
Burst capacity (cache)	NA	500 TB	3 PB	5 PB
Mid-tier capacity (disk)	20 PB	100 PB	1 EB	5 EB
Bottom-tier capacity (tape)	100 PB	1 EB	10 EB	50 EB
I/O servers	400	500	600	700

Implies continuous change!

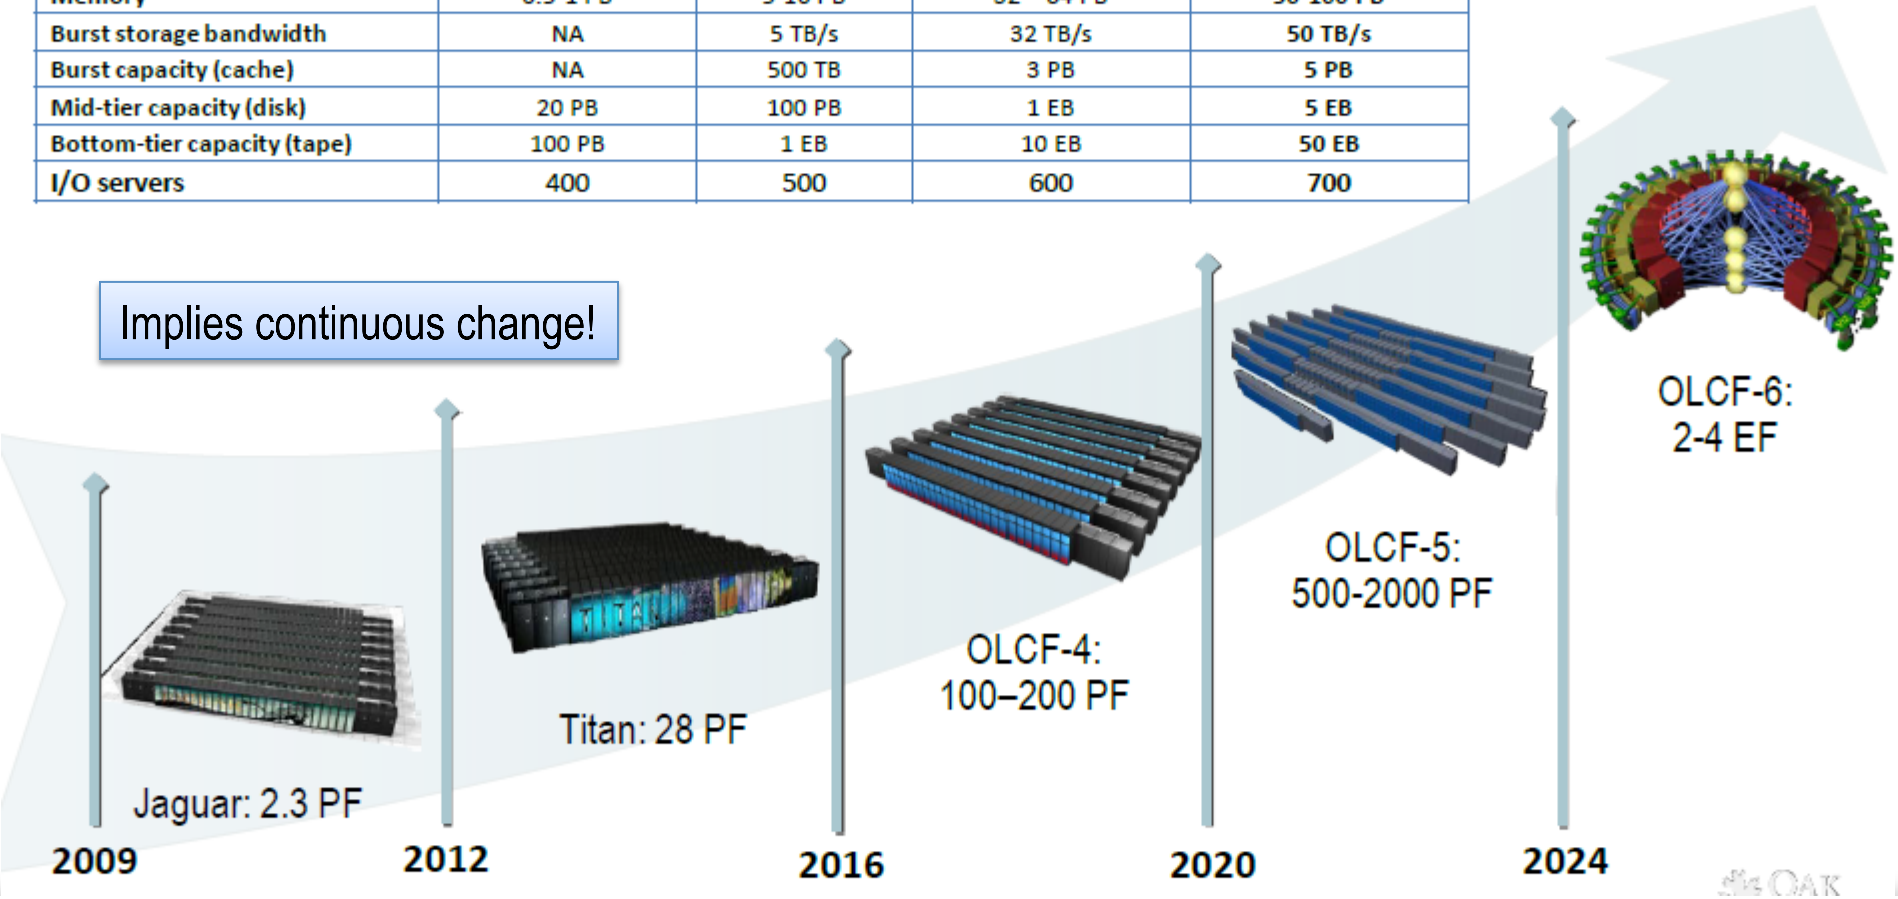

Slide courtesy of W. Joubert

HEP computing report: opportunities

1. Code modernization, maintenance, and dissemination
2. Common tools and coding standards; reduced software footprint
3. Resource support models for smaller-scale projects
4. Data preservation policy for HEP community
5. Distributed Center for Computational Excellence
6. Multi-level computer and computational science training program
7. Community-based expert group for HEP computing
8. Expansion of current interactions with researchers in external disciplines, particularly those in DOE-ASCR

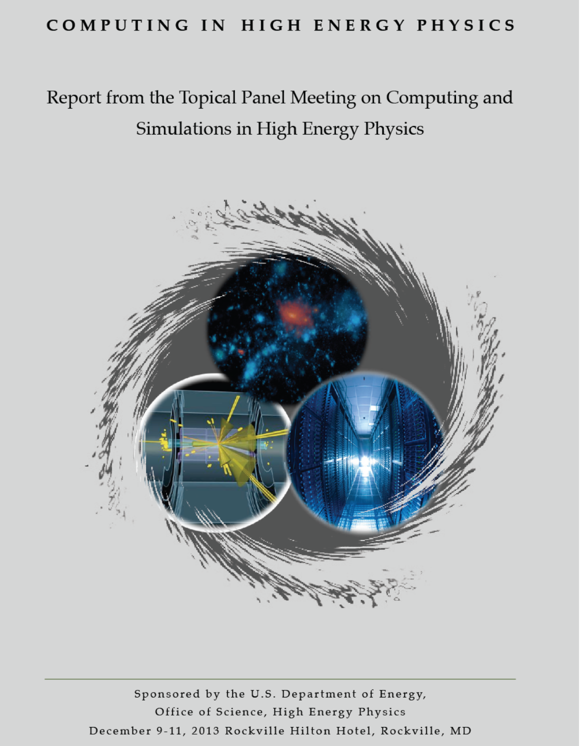

http://science.energy.gov/~/media/hep/pdf/files/Banner%20PDFs/Computing_Meeting_Report_final.pdf

	HEP FY2014 ($M)	NP FY2014 ($M)
Theory	51.2	38
Computation (SciDAC, +)	8.5	2.0
Total	797	489
Ratios	6% (Th/Tot) 1% (Comp/Tot)	7% (Th/Tot) 0.5% (Comp/Tot)

Materials Genome…

BES opportunities

Vision: *Advanced materials are essential to economic security and human well-being … the Materials Genome Initiative* ***will enable discovery, development, manufacturing, and deployment of advanced materials at least twice as fast as possible today, at a fraction of the cost.***

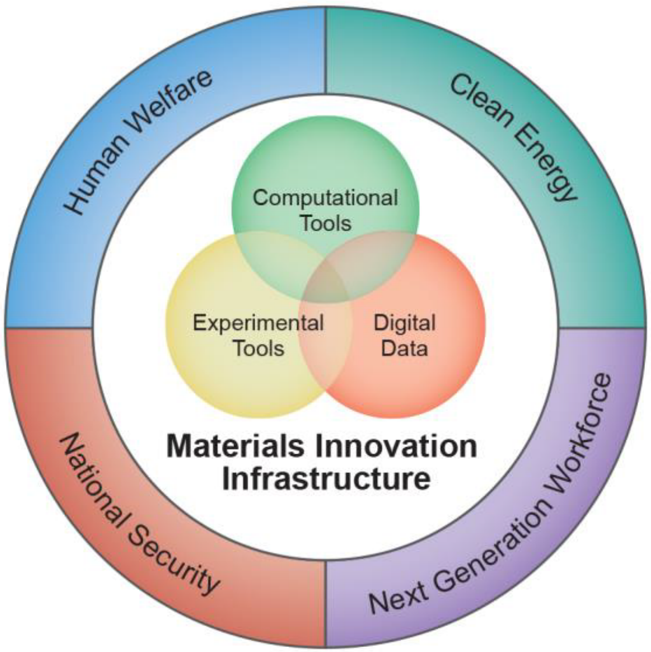

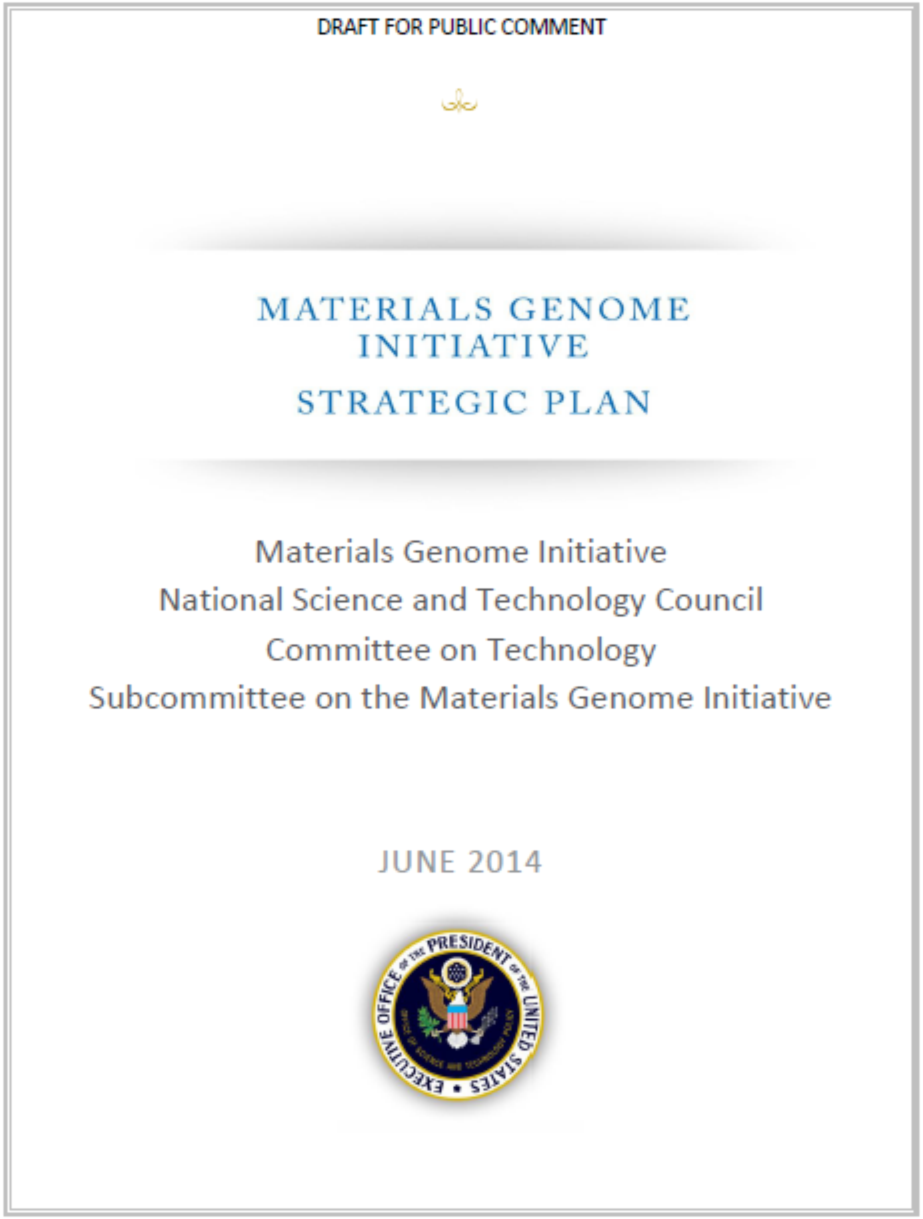
DRAFT FOR PUBLIC COMMENT

MATERIALS GENOME INITIATIVE

STRATEGIC PLAN

Materials Genome Initiative
National Science and Technology Council
Committee on Technology
Subcommittee on the Materials Genome Initiative

JUNE 2014

Four challenges:

(1) Leading **a culture shift in materials research** to encourage and facilitate an integrated team approach that **links computation, data, and experiment** and crosses boundaries from academia to industry;

(2) **Integrating experiment, computation, and theory** and **equipping the materials community with the advanced tools and techniques to work across materials classes** from research to industrial application;

(3) Making digital data accessible including combining data from experiment and computation into a searchable materials data infrastructure and encouraging researchers to make their data available to others;

(4) Creating a world-class materials workforce that is trained for careers in academia or industry, including high-tech manufacturing jobs.

Covered by "Computational Materials Sciences" in BES FY2015 PBR: +$25M

Network for ab initio many-body methods: development, education and training

- Principal Investigator: Paul Kent (ORNL)
  - Co-Investigators: David M. Ceperley, University of Illinois
  - Miguel A. Morales, LLNL
  - Jeff Greeley, Purdue University
  - Luke Shulenberger, SNL
- This project links the developers of ab initio many-body electronic structure methods, especially quantum Monte Carlo (QMC) methods, and the developers and users of QMCPACK (an open-source QMC package) to build a next-generation QMC framework that accelerates discovery of advanced materials.
- (i) heterogeneous catalysis of metallic nanoclusters and nanoparticles,
  (ii) defect formation, energetics and effects on materials properties, and
  (iii) phase transitions and properties of materials under pressure.
- QMC is one of the very few electronic structure methods that have the potential to produce systematically improvable predictions for condensed matter systems. The results from this project to date demonstrate that while significant challenges remain, many complex materials are within scope of the current methods, algorithms and computational resources.

Developing an LRP initiative

- Establish the need – why do we compute?
- Are we competitive with HEP and BES?
  - Answer: no…we are behind the curve
  - We could be relying too much on the SciDAC model
- What about experimental data?
  - LHC/ALICE computational needs through NERSC and ORNL
  - RHIC/TJ data? EIC data? FRIB data?
- Crucial to tie to experimental efforts: $0\nu\beta\beta$, FRIB, TJ, RHIC, LHC…
- Any large scale investment will be long-term
  - People, algorithms, and maybe hardware
  - Cannot depend on ASCR to fix *all* our problems
  - How would these investments be used to fill gaps (and what are the gaps)?
- A statement like 'we need computing', is NOT sufficient – we need a specific recommendation

In order to accelerate discovery in nuclear physics, and to bridge the gap between the highest-end computational platforms provided by ASCR and conventional single-investigator platforms, the Office of Nuclear Physics should invest $XX M/year in computational infrastructure (people and hardware?). [Follow this by a paragraph that details why this is needed and why now]