FUTURE TRENDS IN
NUCLEAR PHYSICS
COMPUTING

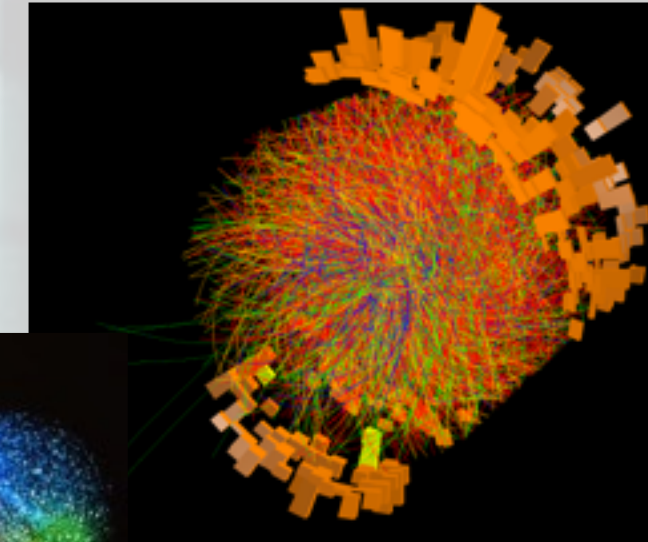

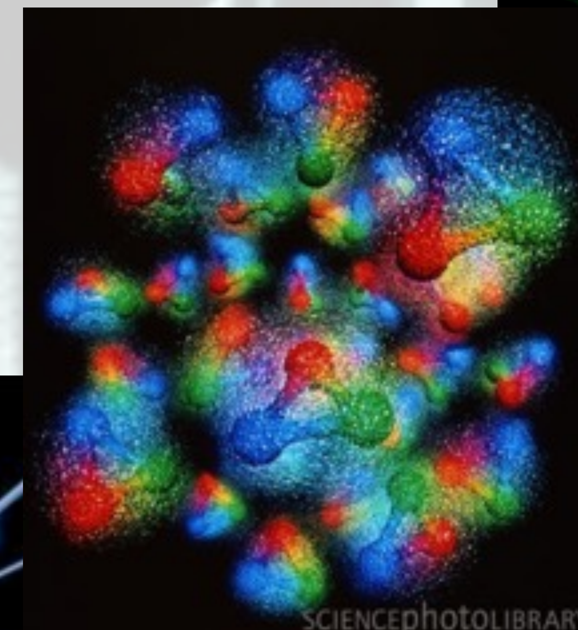

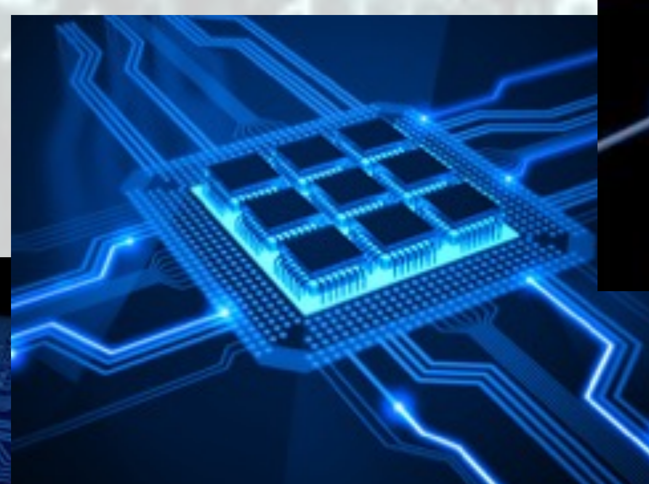

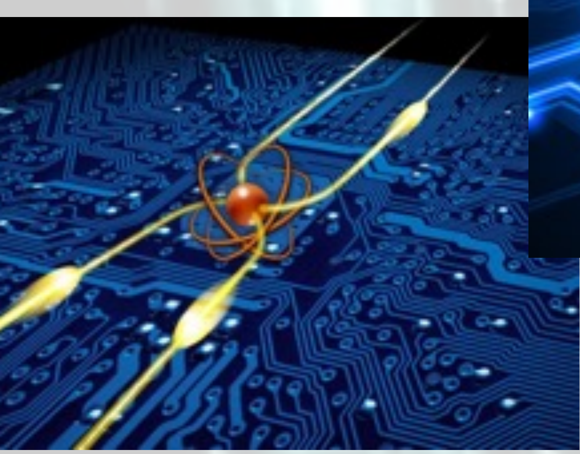

Nuclear Physics Computing in a Decade

Future Trends in Nuclear Physics Computing
Jefferson Laboratory, May 2, 2017

Martin J Savage

Institute for Nuclear Theory
University of Washington

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

The Next Decade of Scientific Computing will Forever Change Our Understanding of Matter

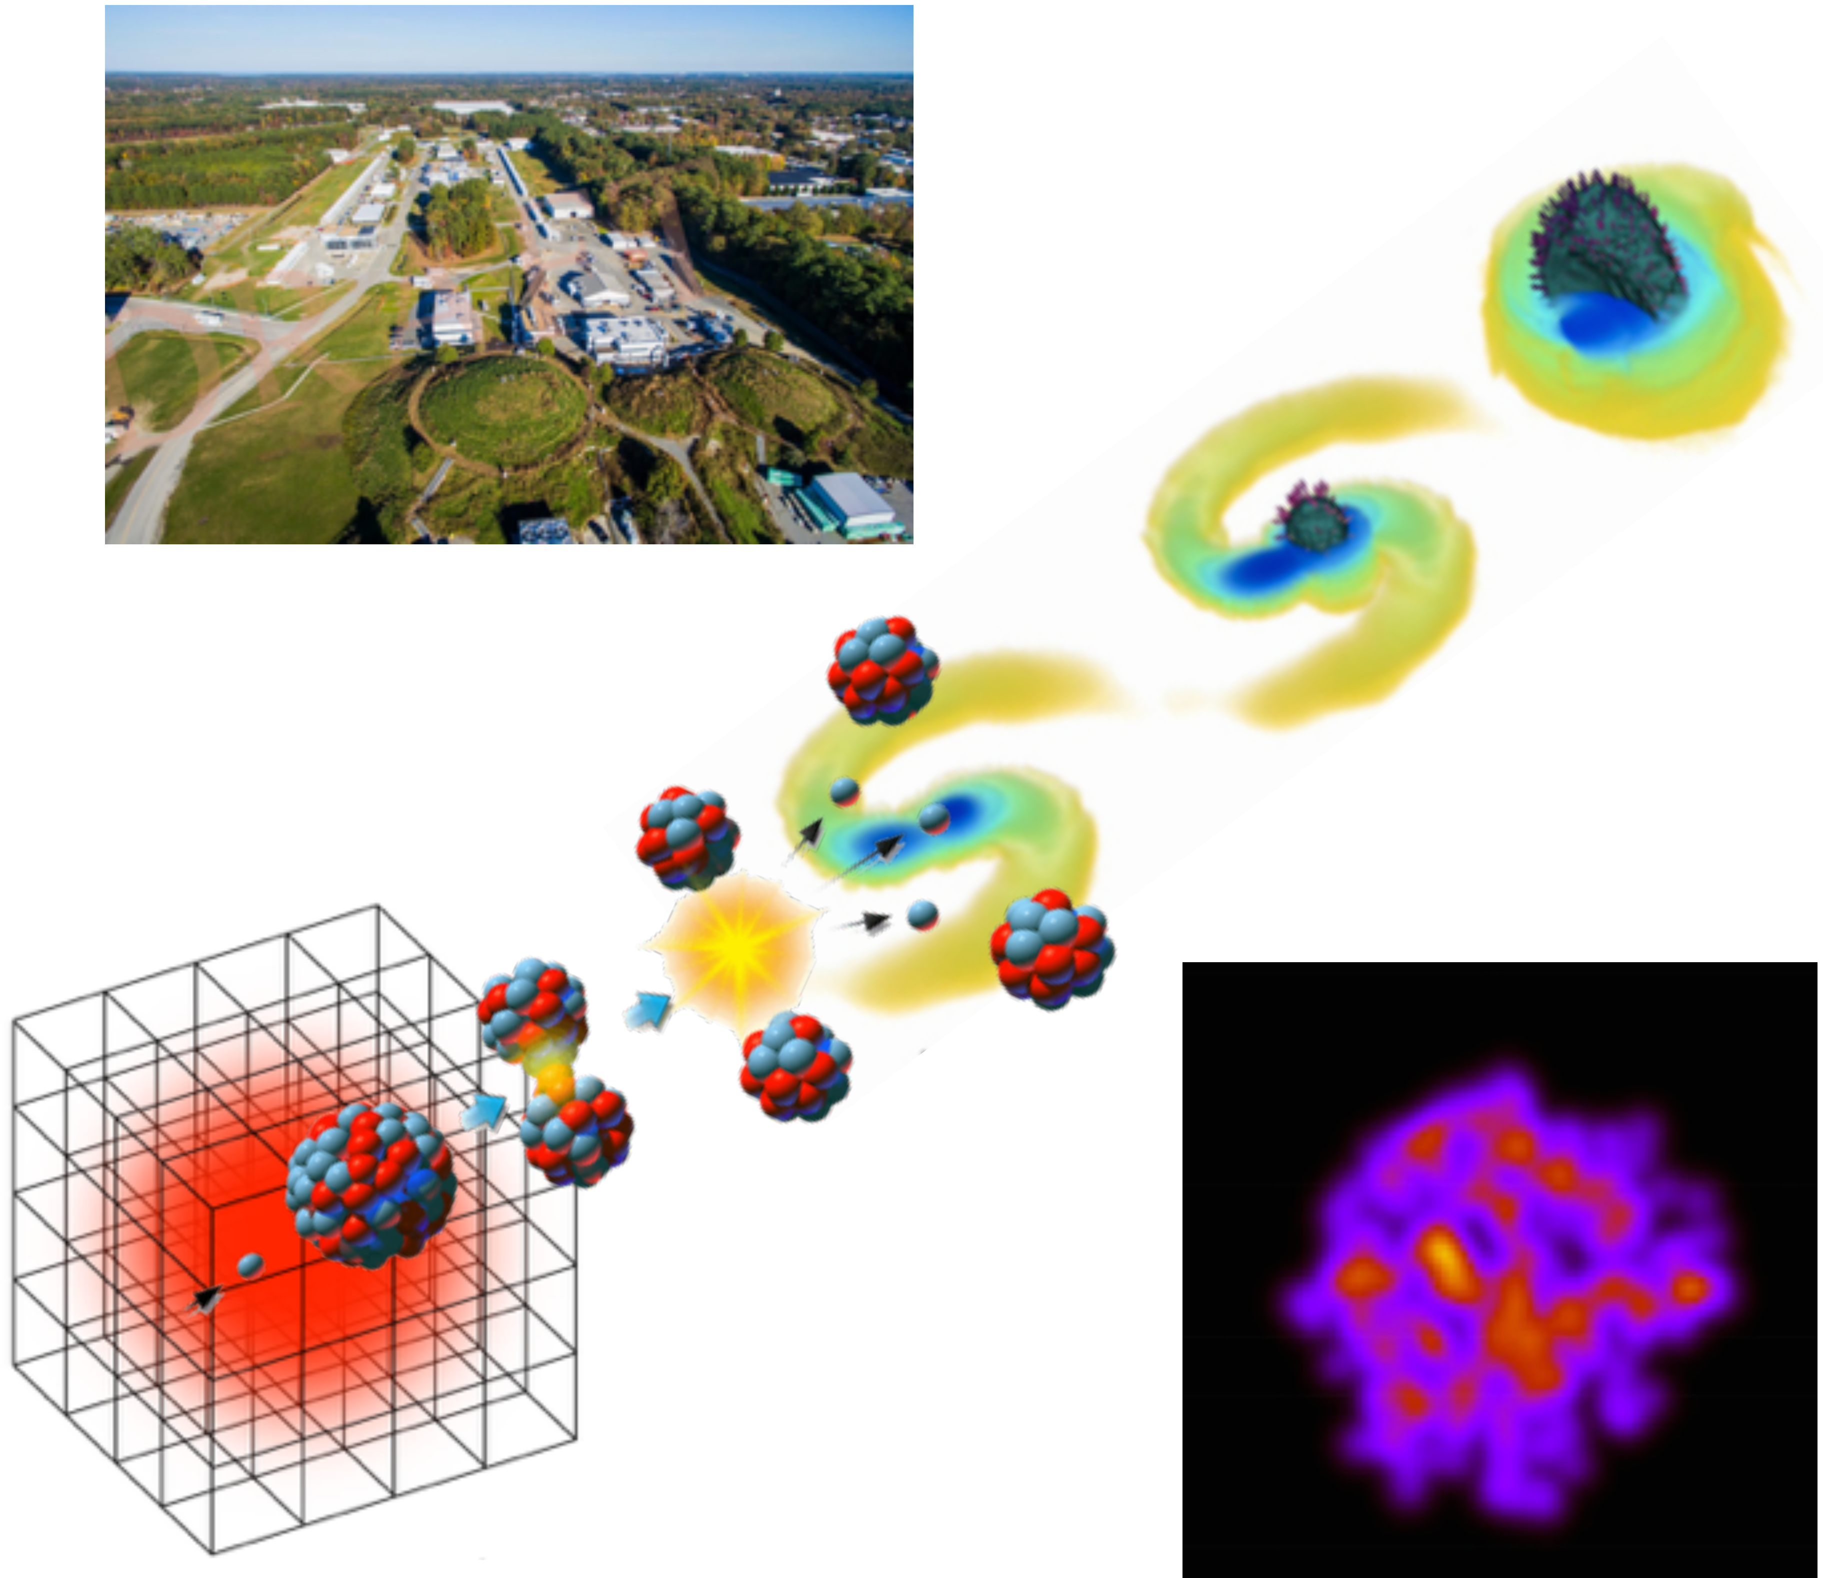

Imagine being able to predict — with unprecedented accuracy and precision — the structure of the proton and neutron, and the forces between them, directly from the dynamics of quarks and gluons, and then using this information in calculations of the structure and reactions of atomic nuclei and of the properties of dense neutron stars...

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Scientific Objectives and Applications Rely on High-Performance Computing

HPC is *essential* in:


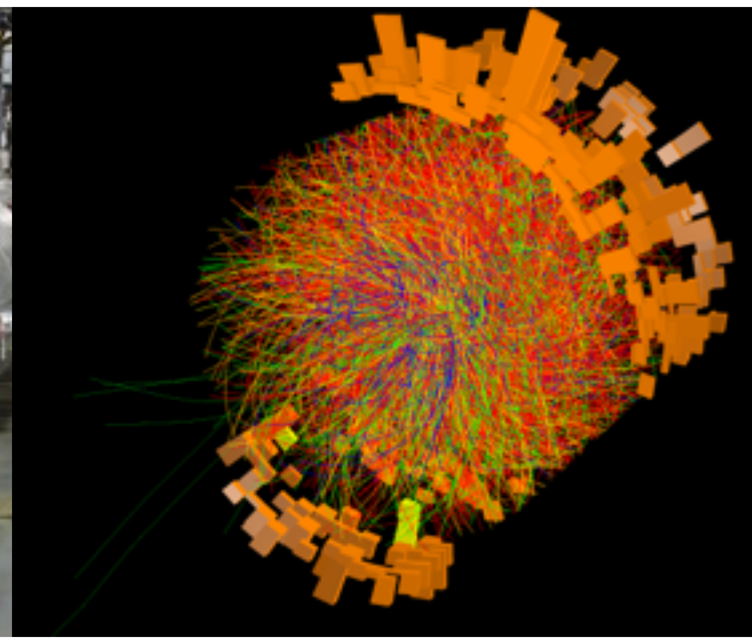

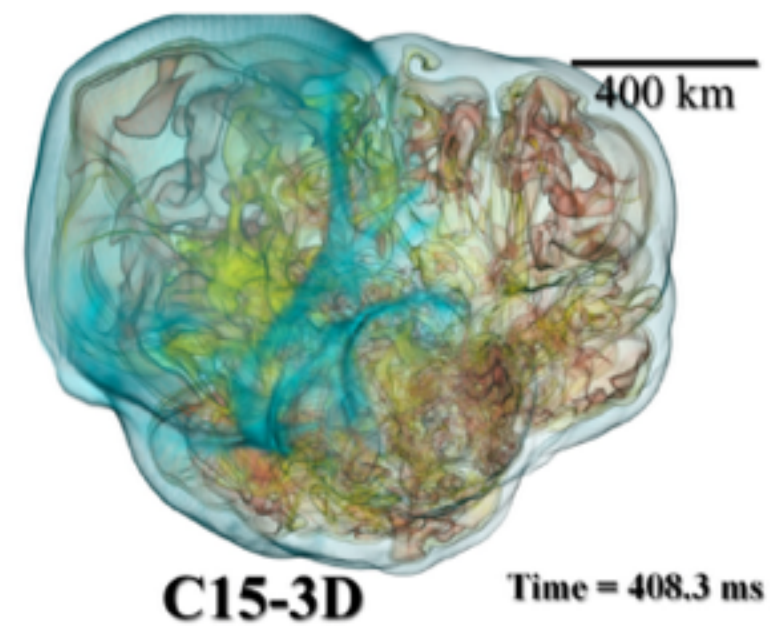

1. Design and Optimization of the extensive and vibrant NP experimental program

2. Acquisition and handling of experimental data

3. Large-scale simulations and calculations of emergent complex systems from subatomic to cosmological

Nuclear Physics Scientific Objectives and Applications Rely on High-Performance Computing

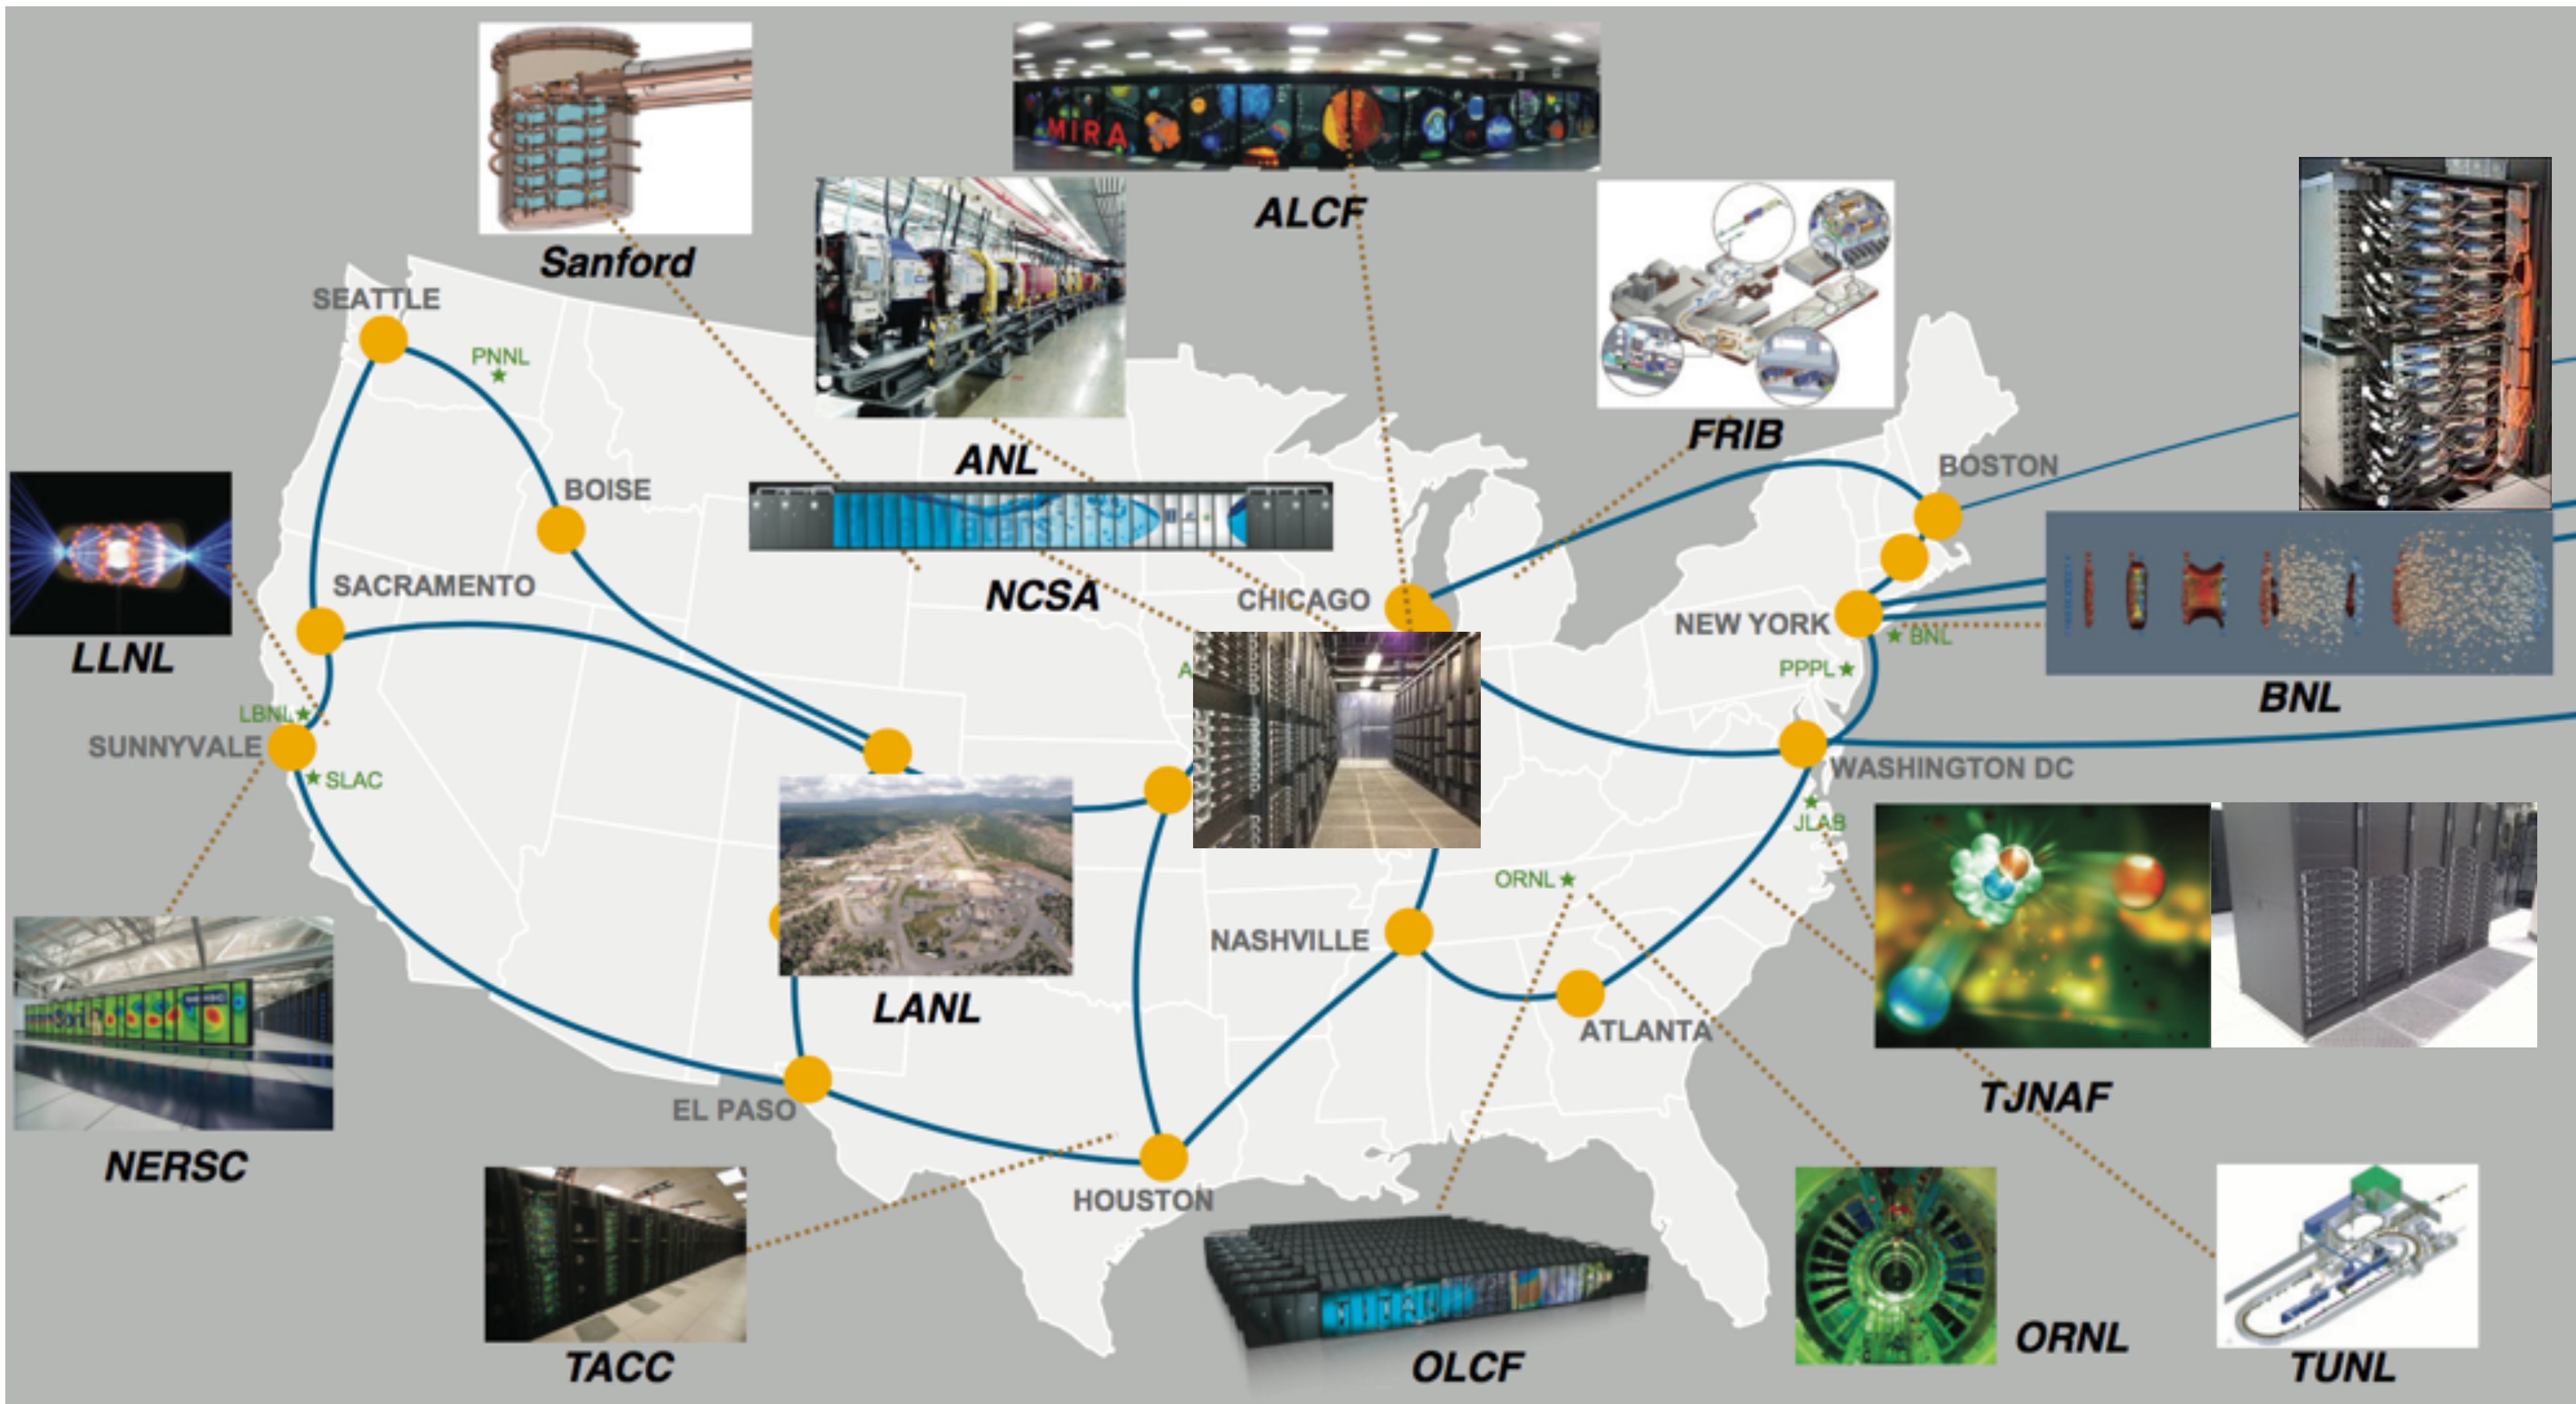

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Preparing for Exascale Computing and Beyond

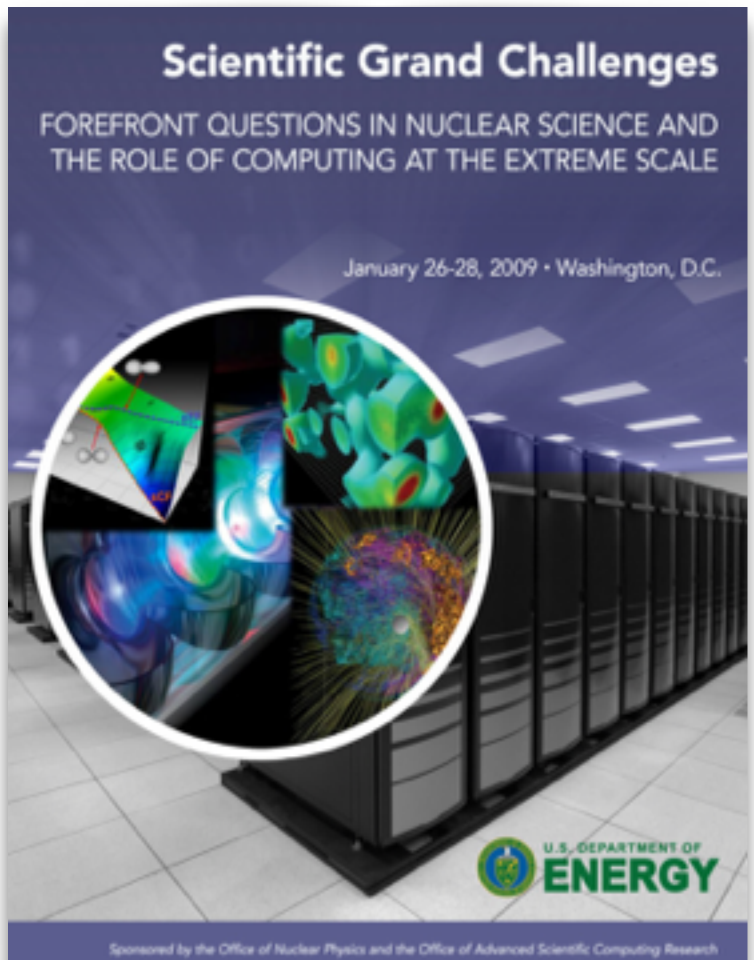

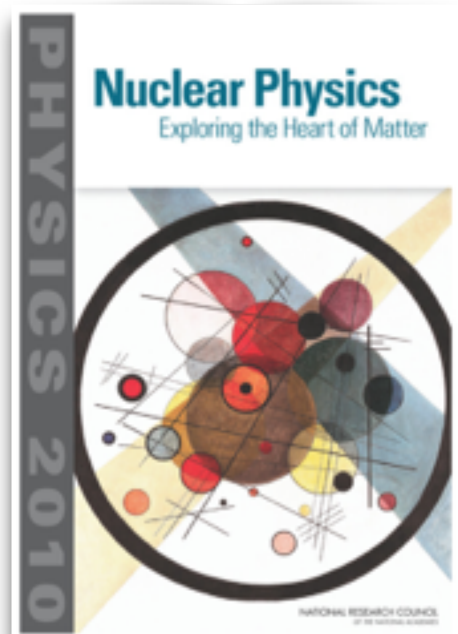

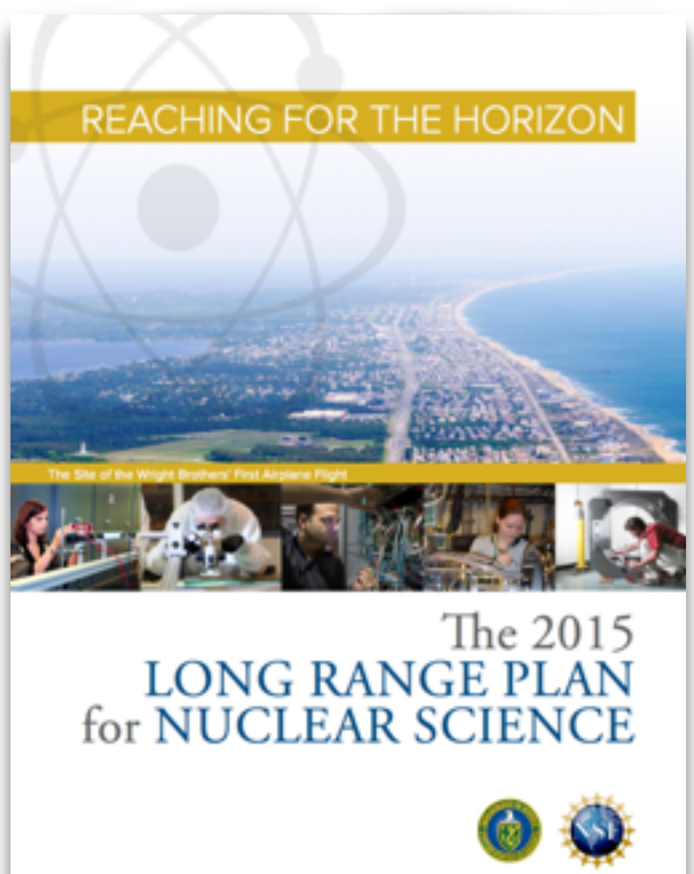

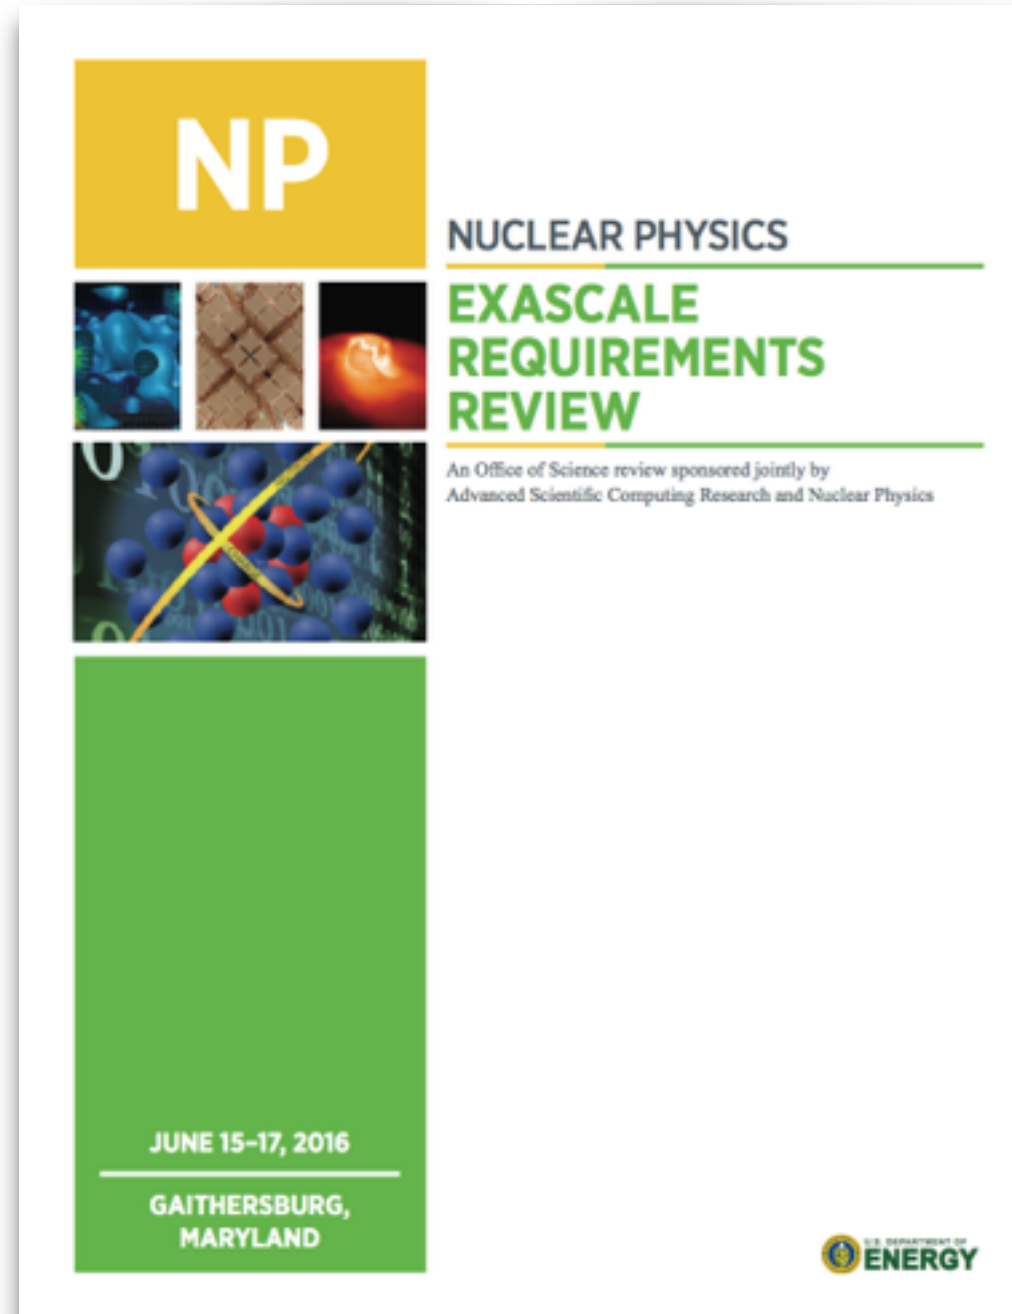

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

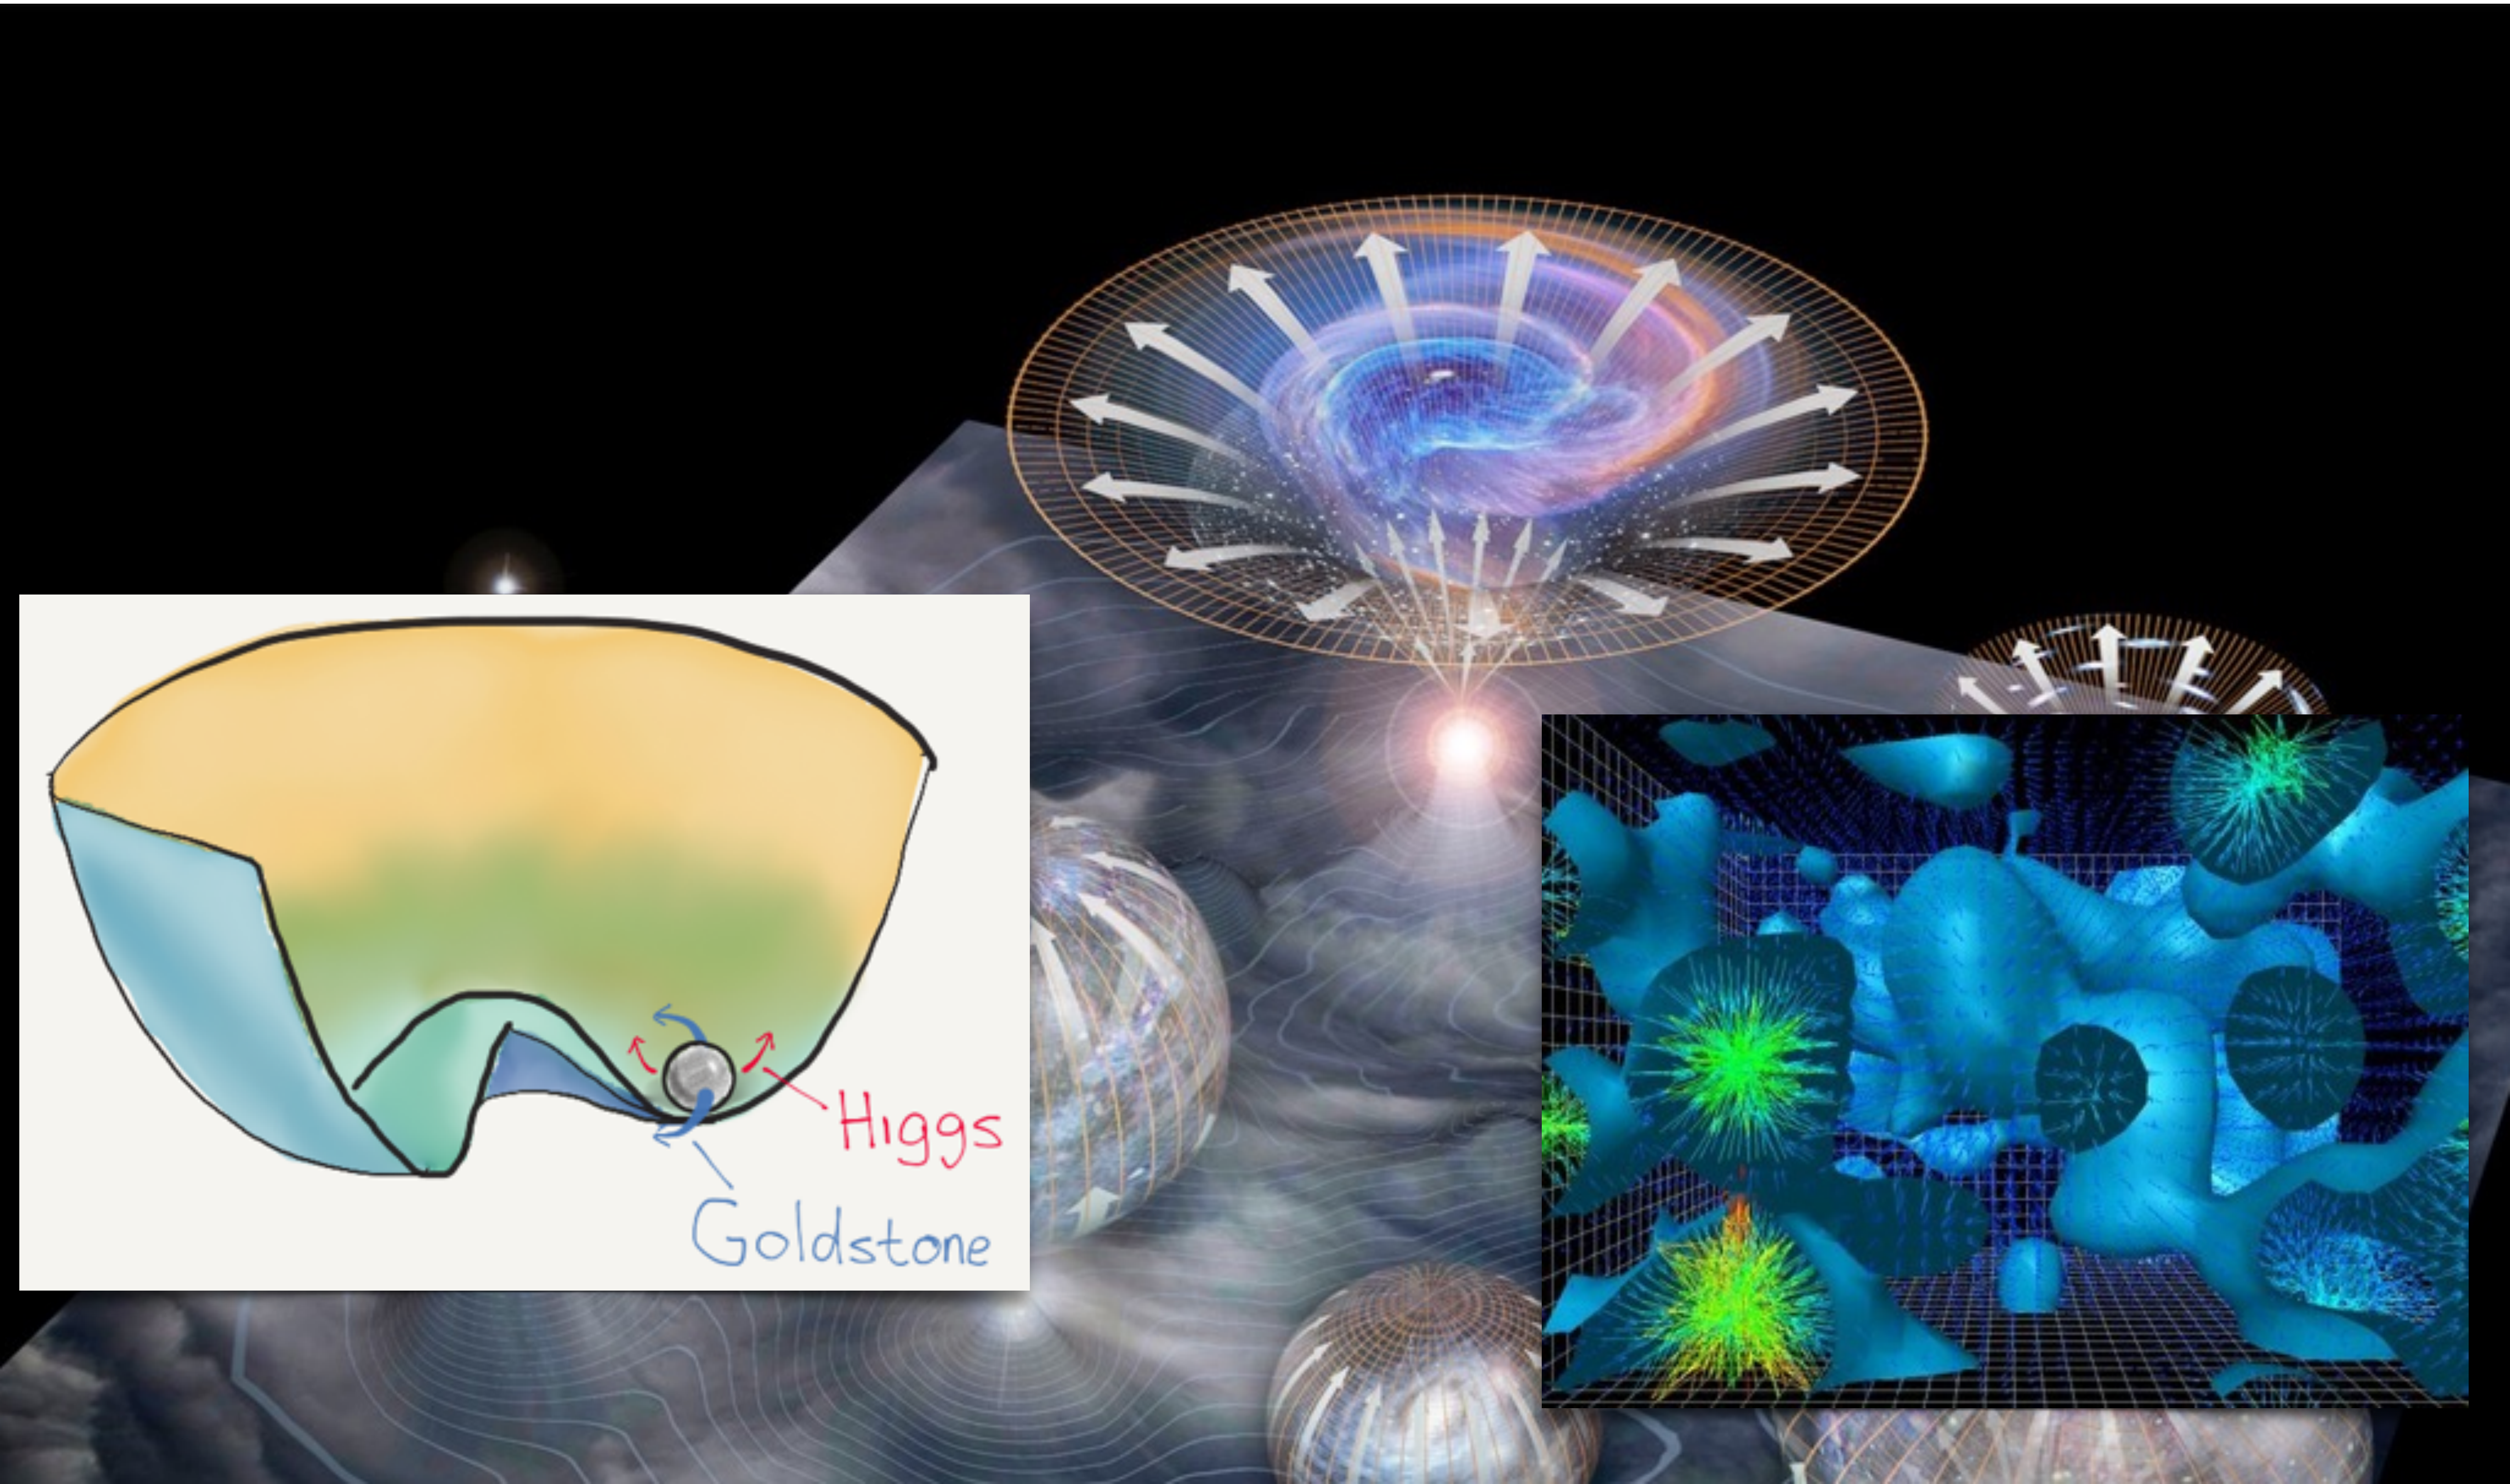

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

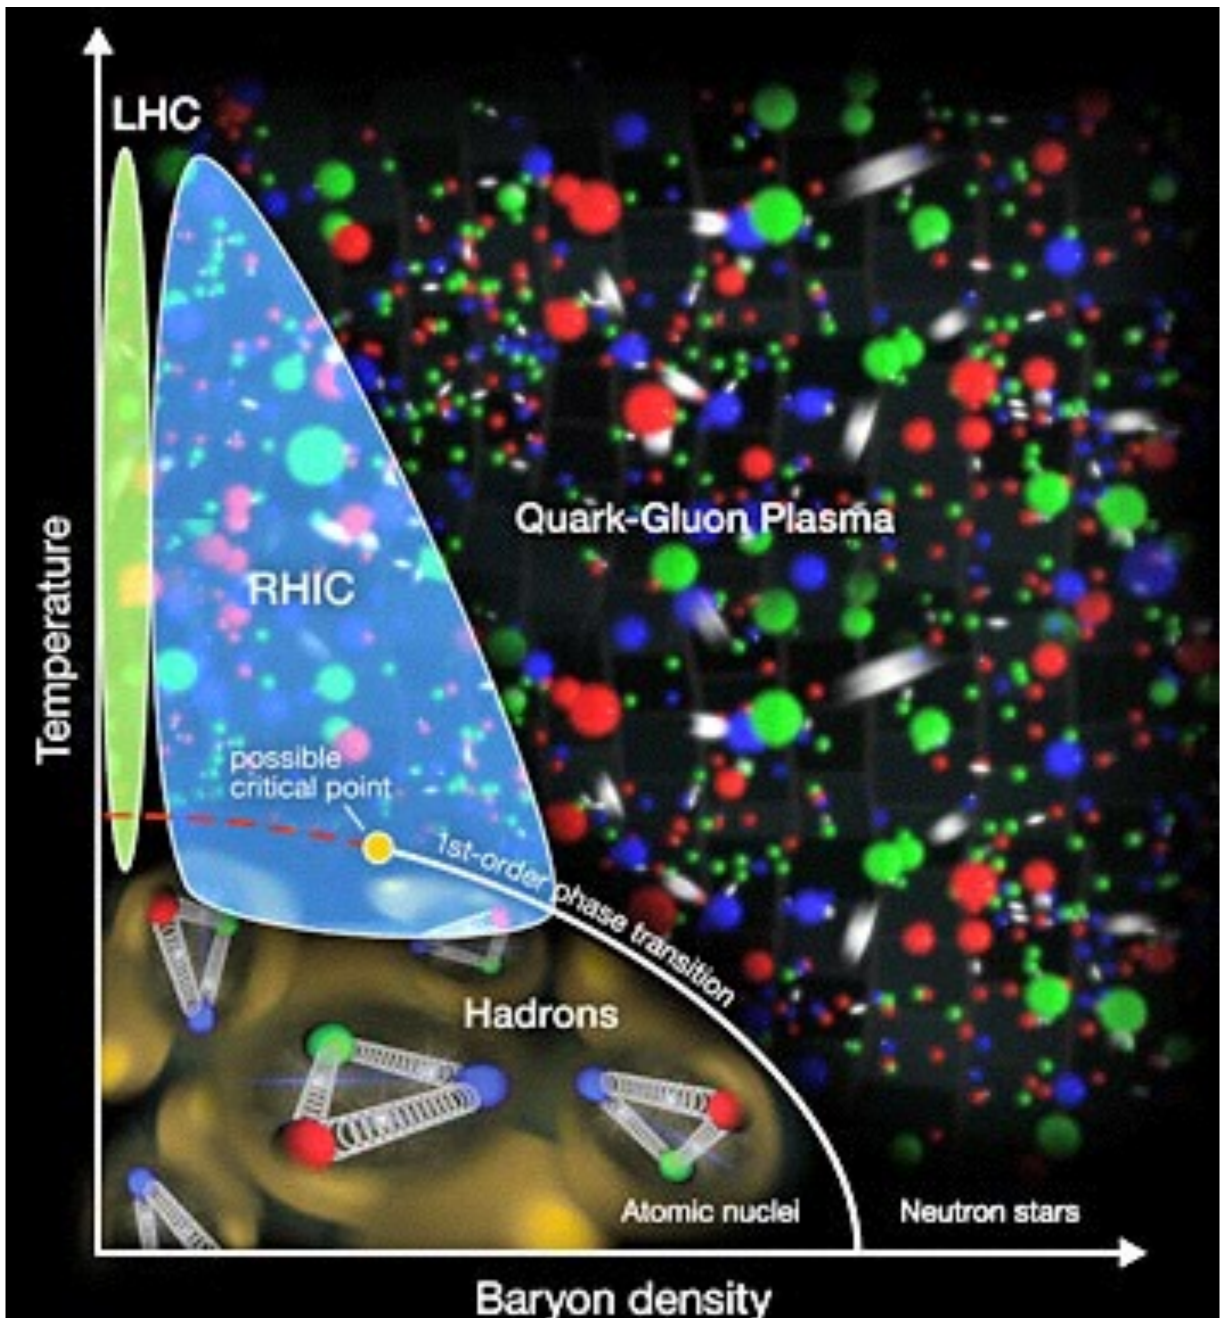

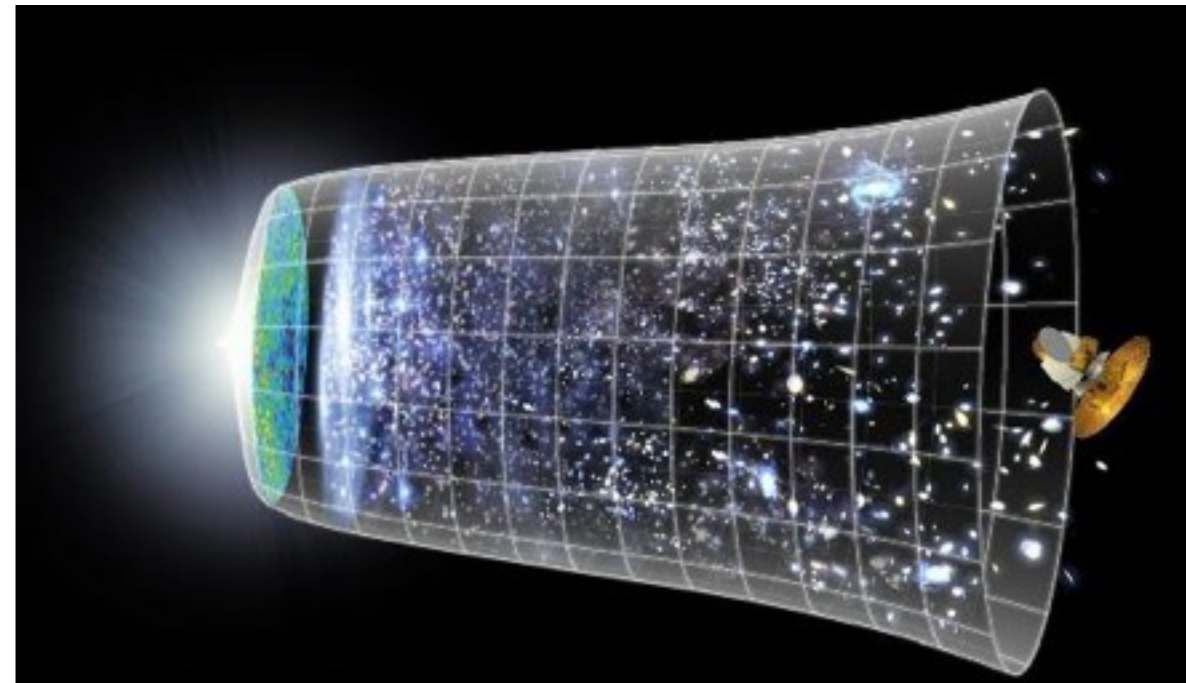

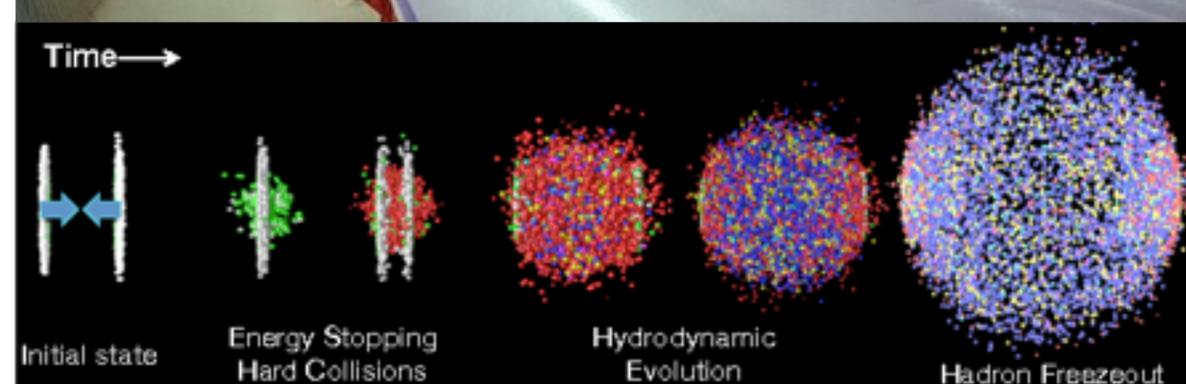

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

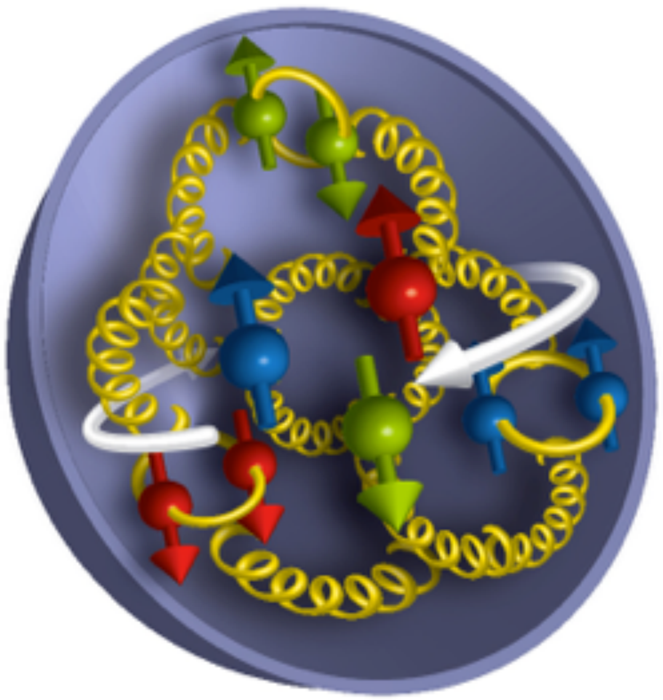

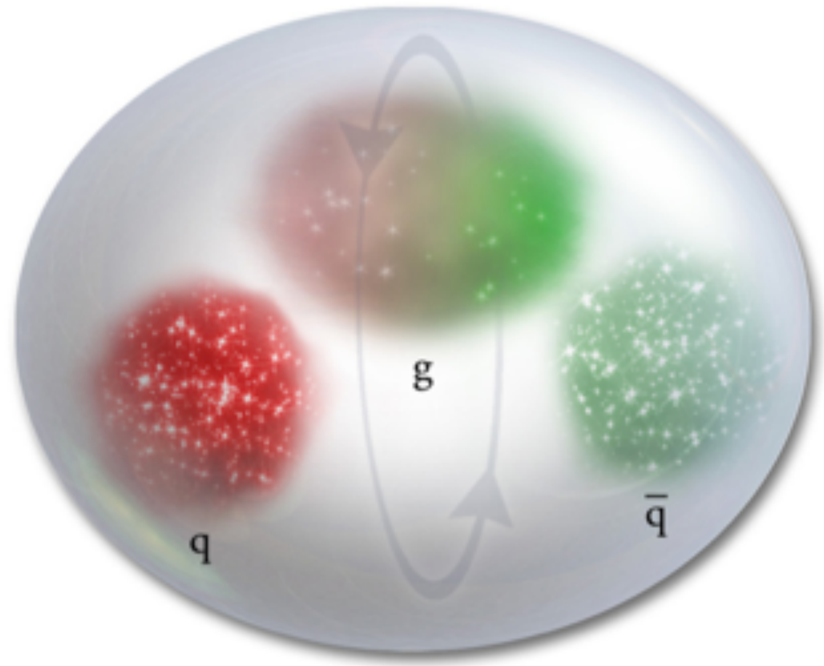

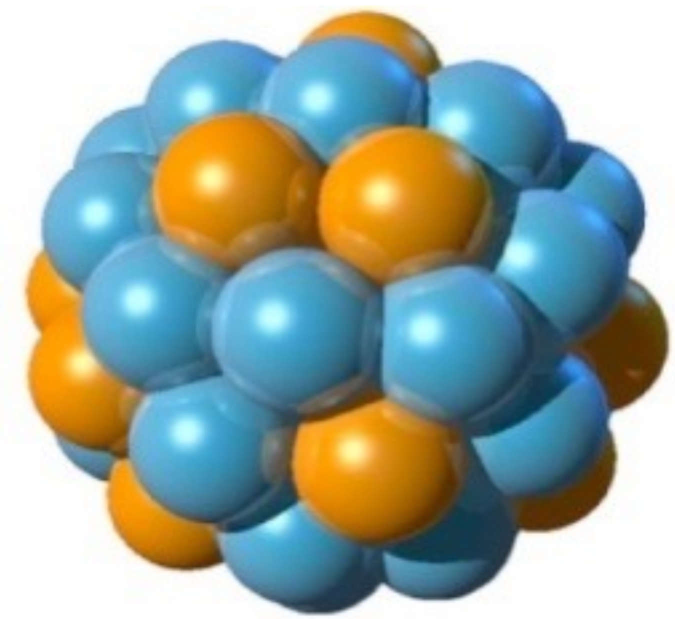

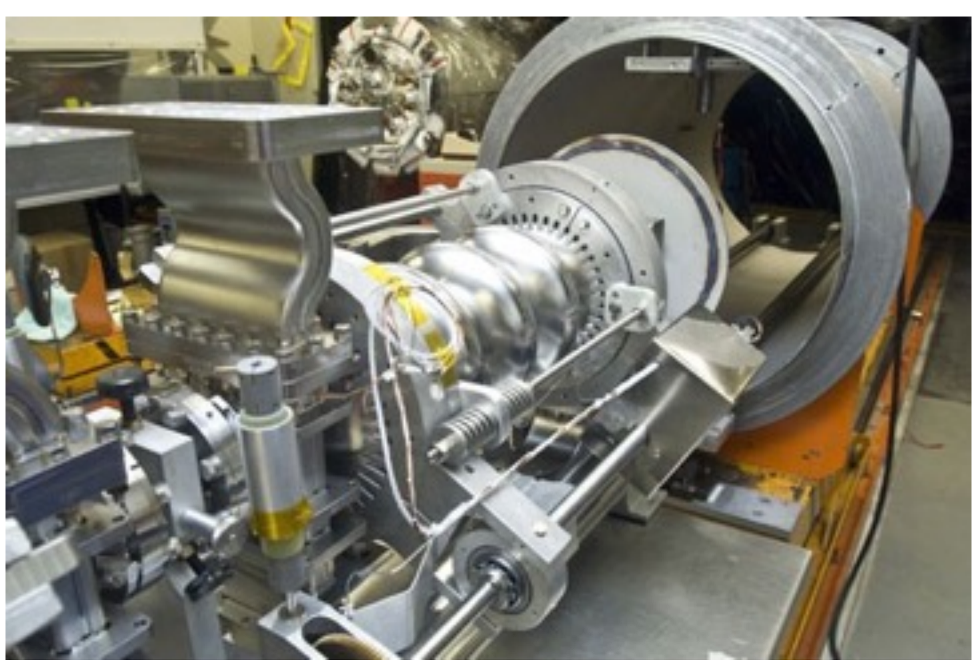

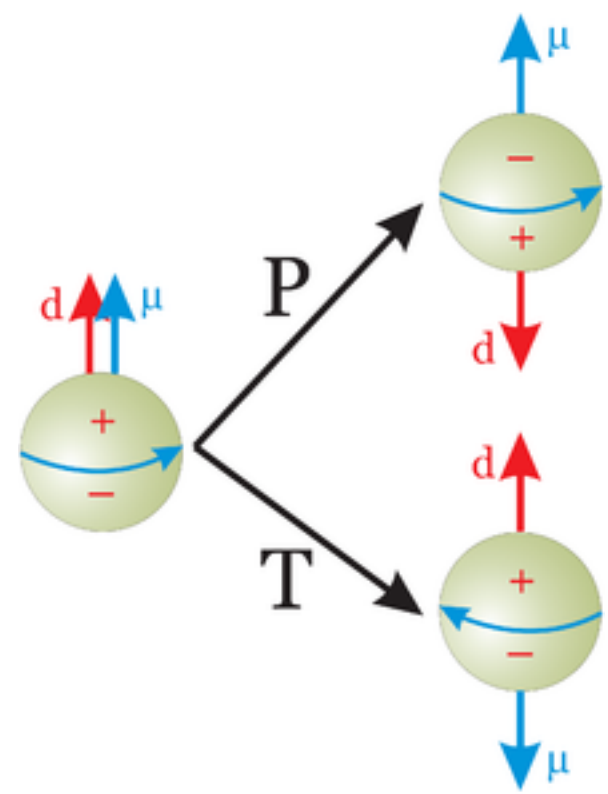

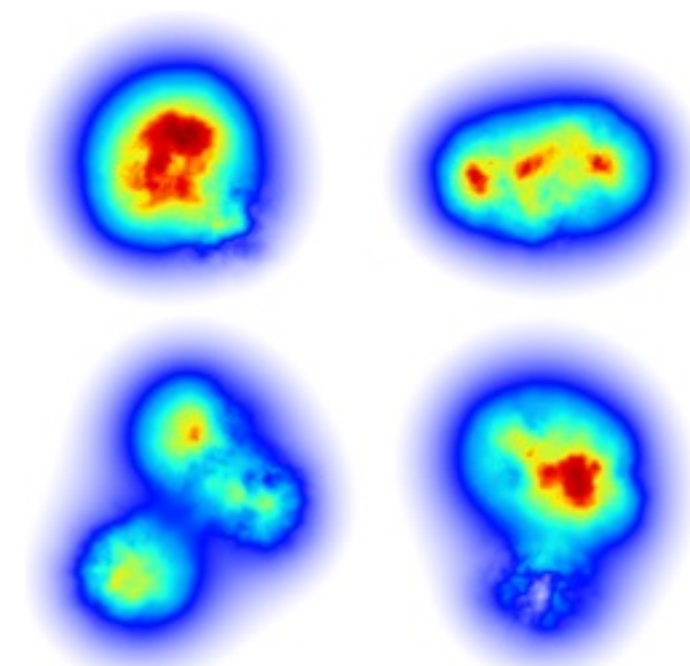

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

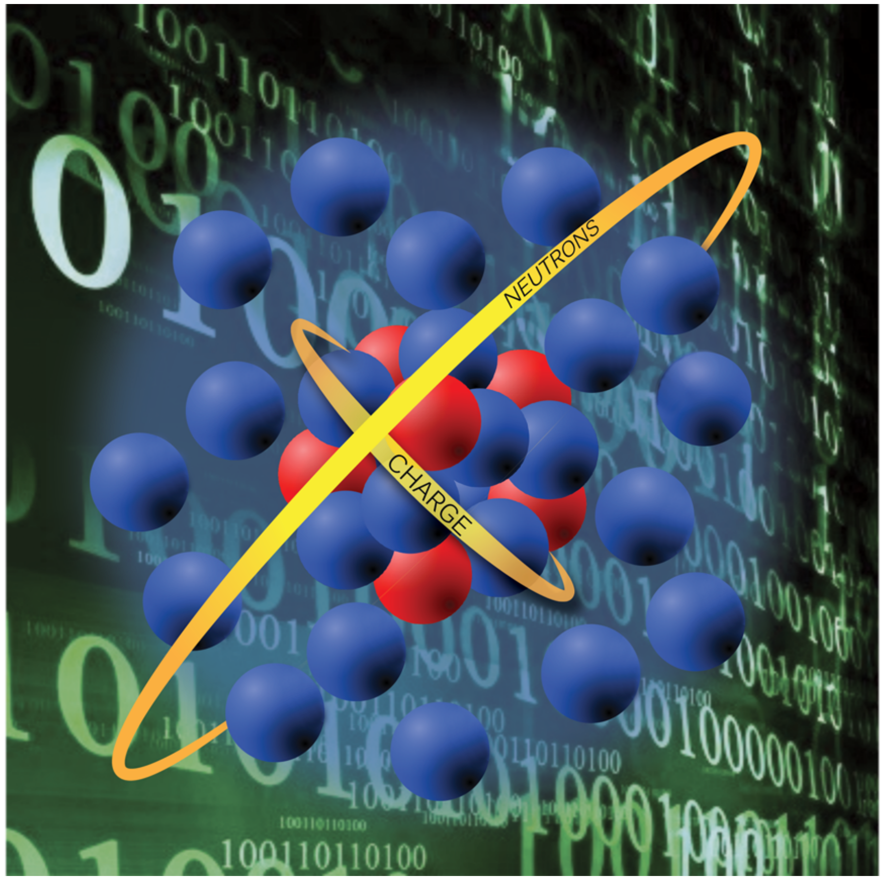

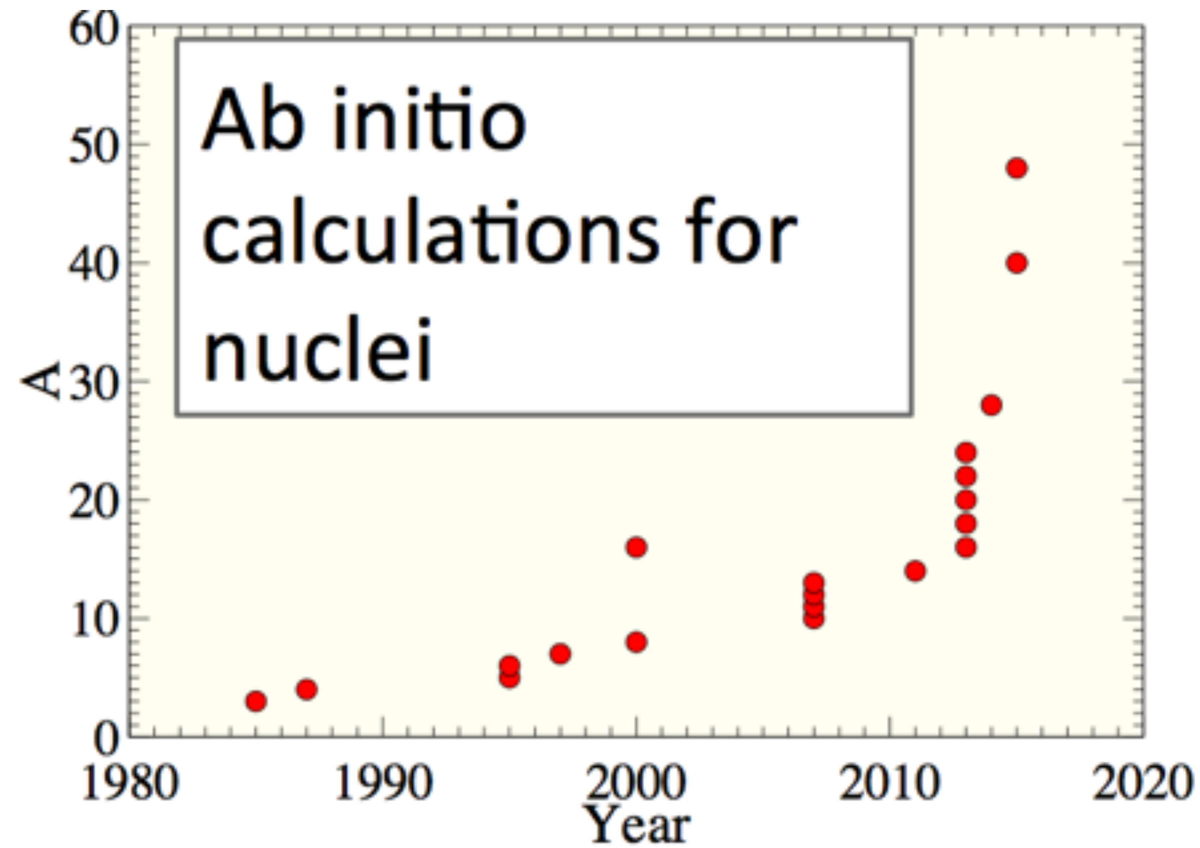

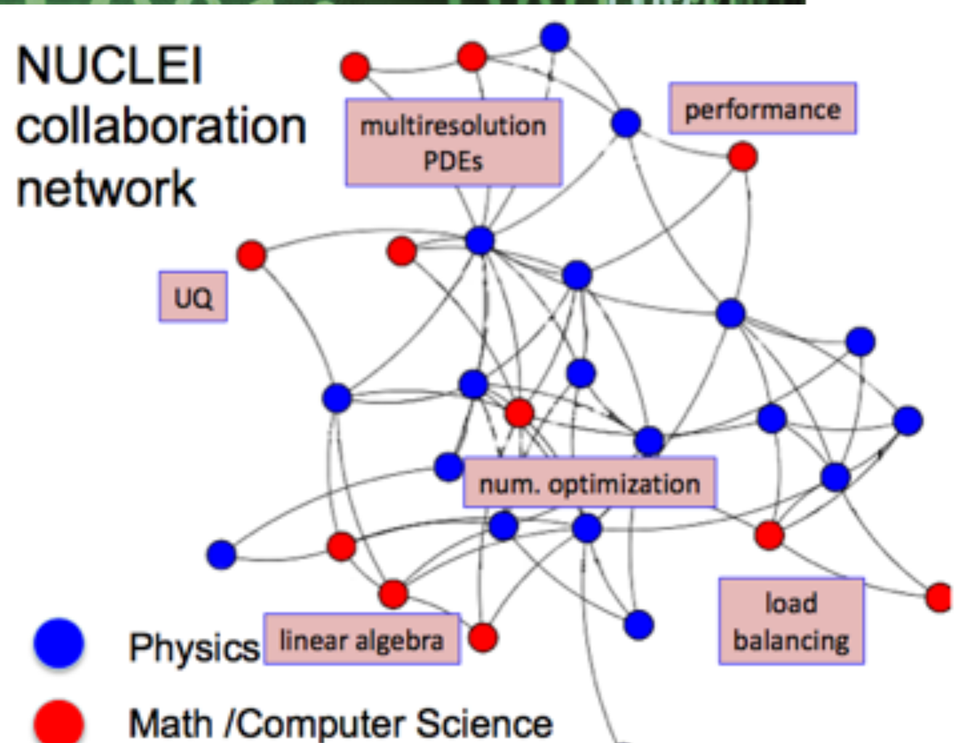

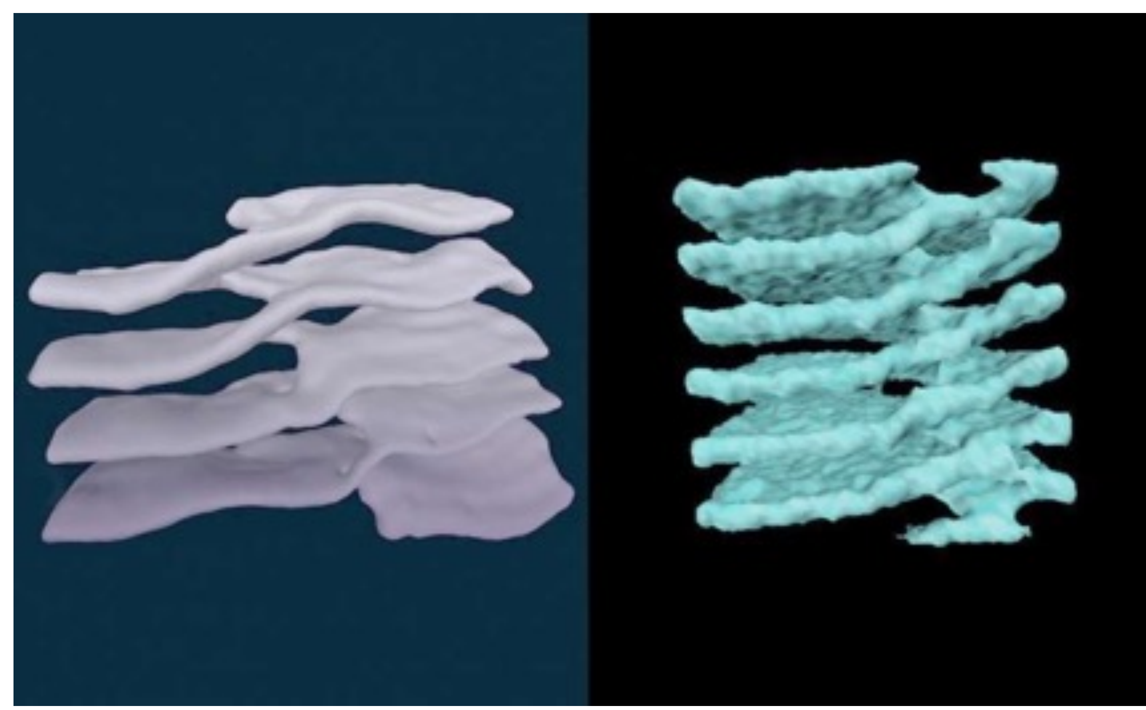

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

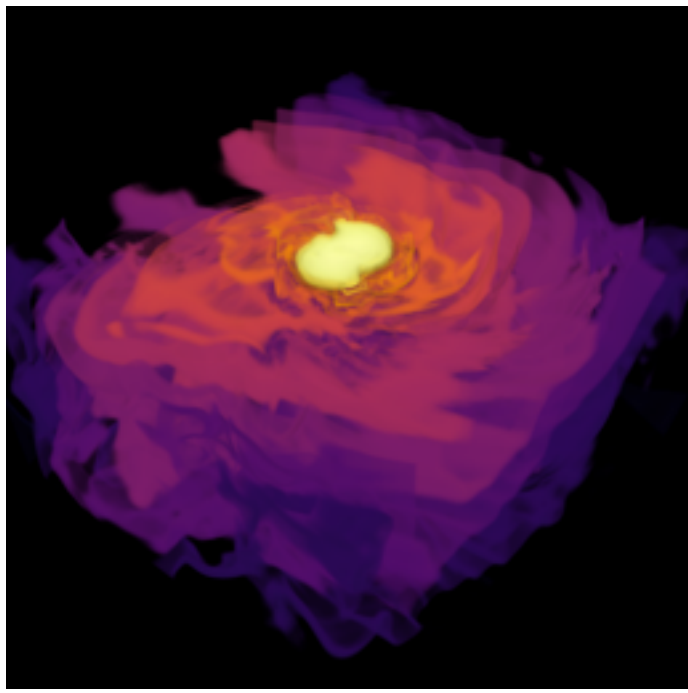

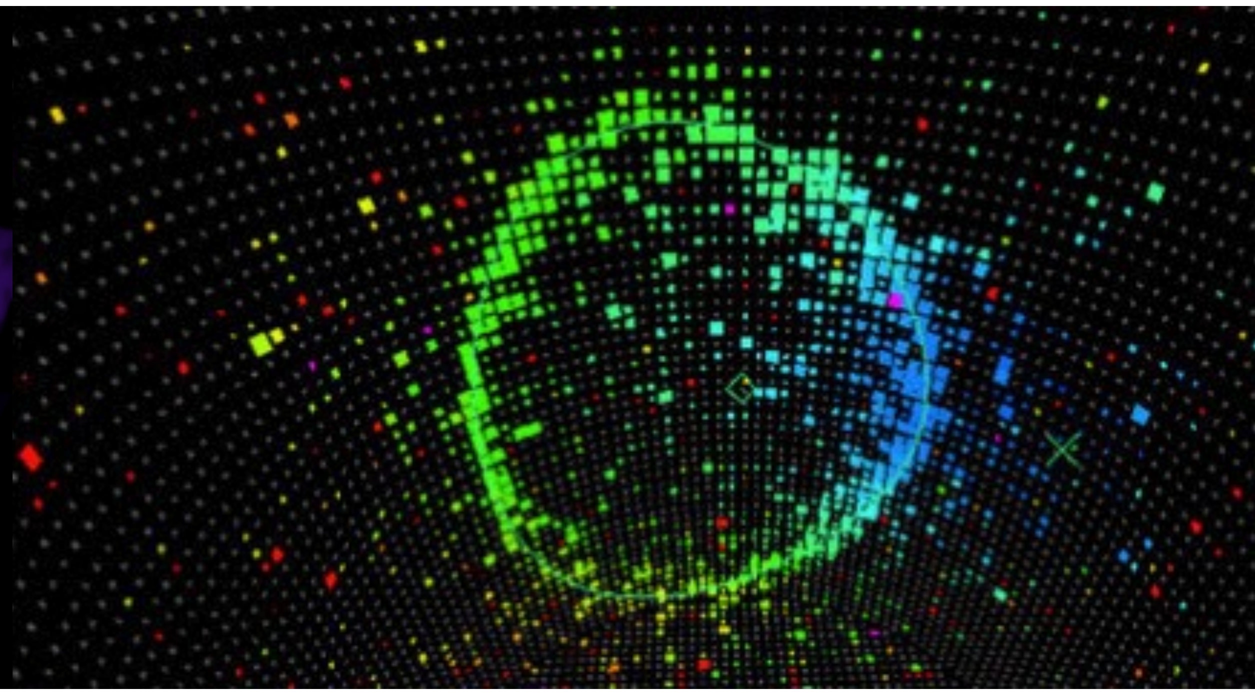

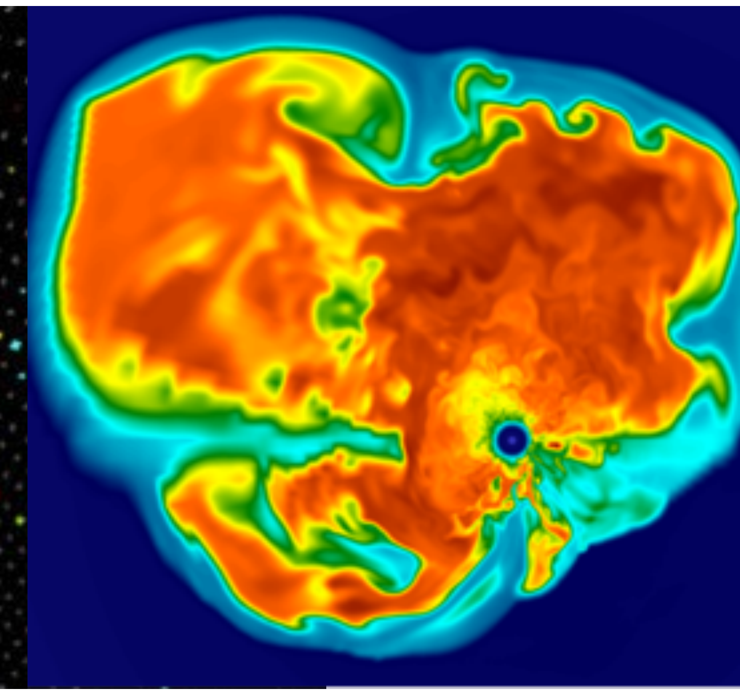

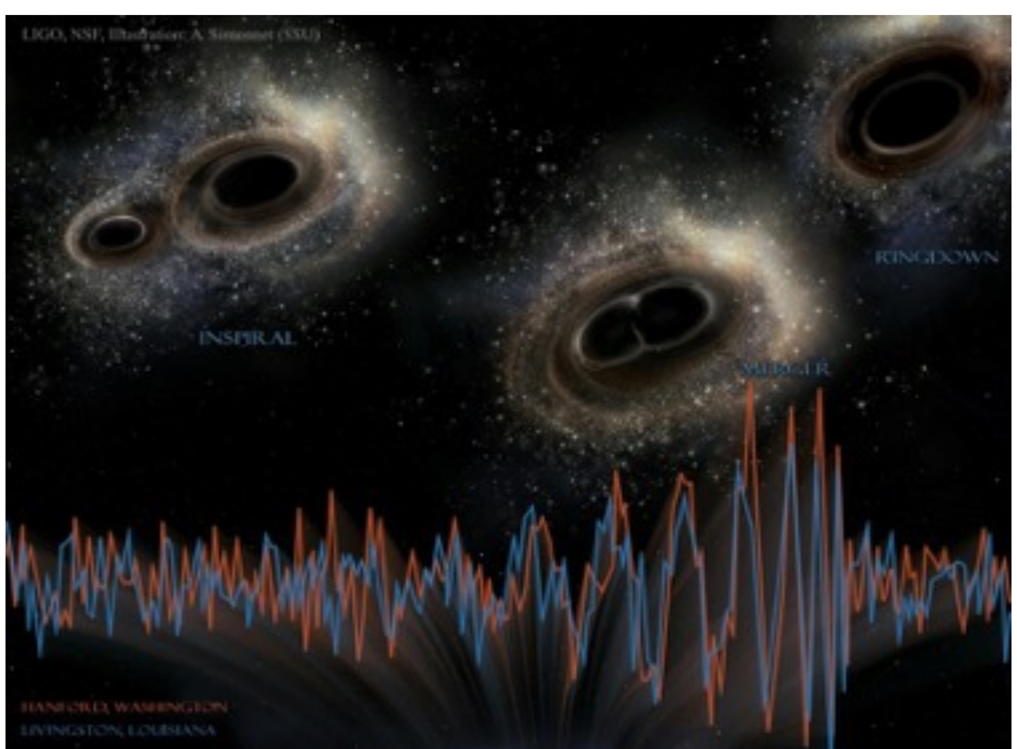

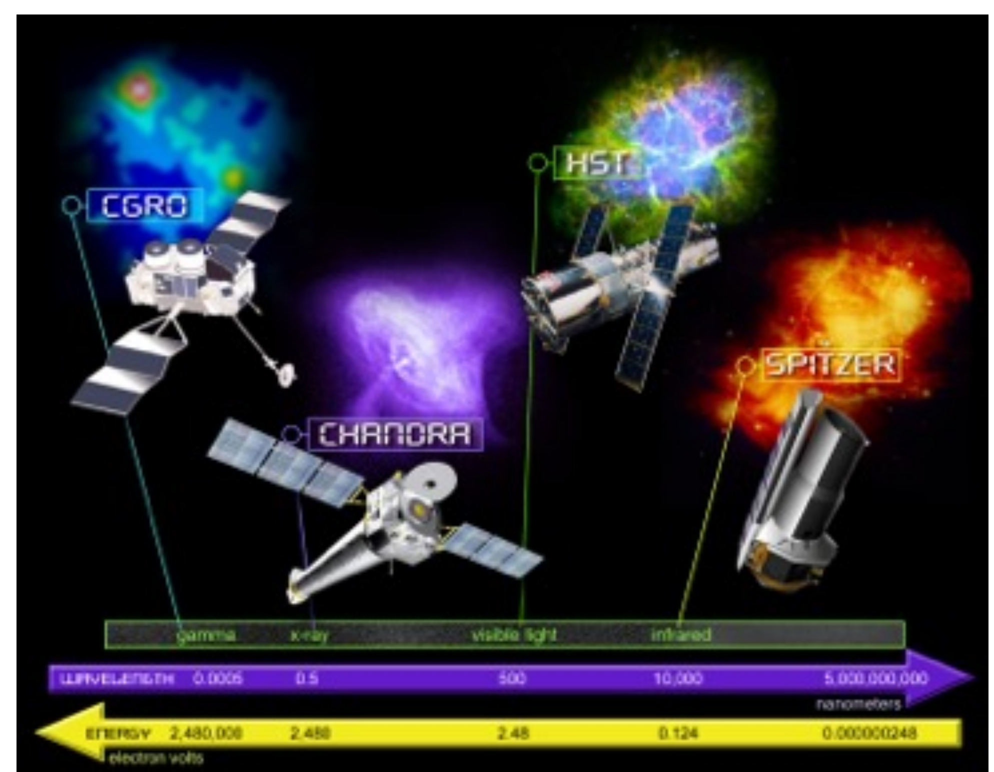

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Nuclear Physics Emerges from Nature's Choice of Fields, their Nonlinear Interactions and Plancks Constant

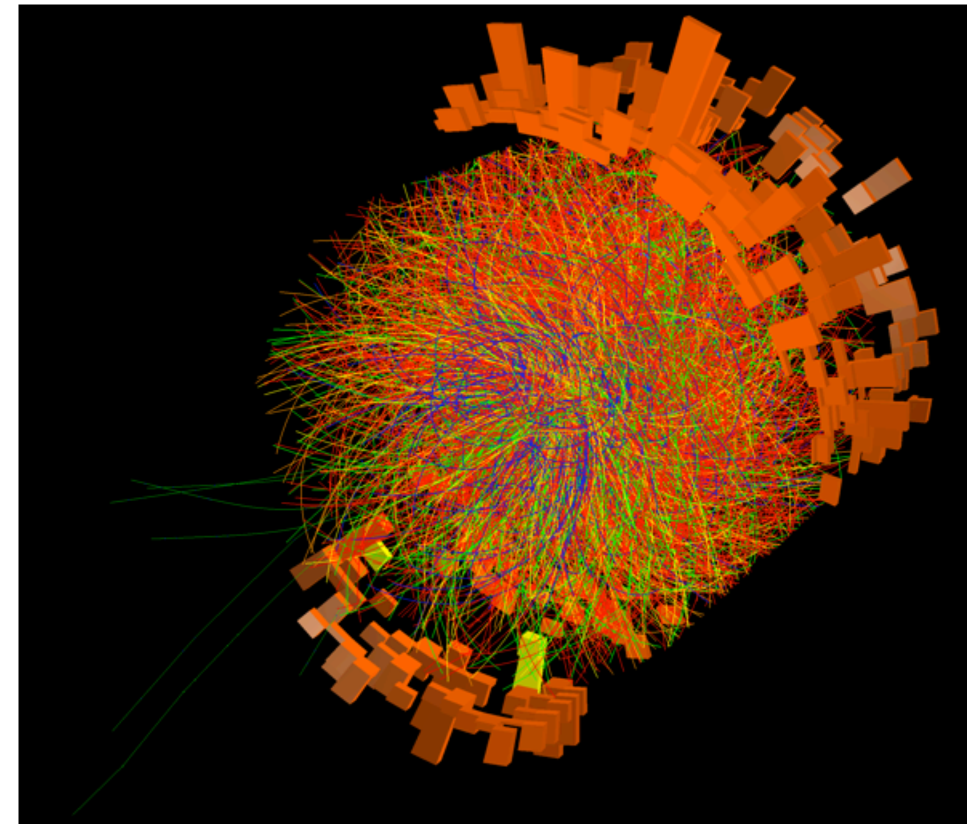

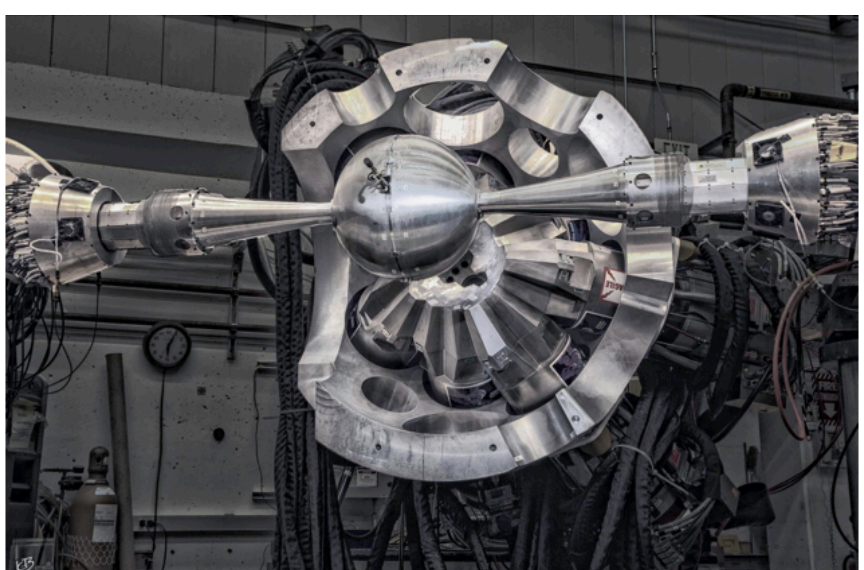

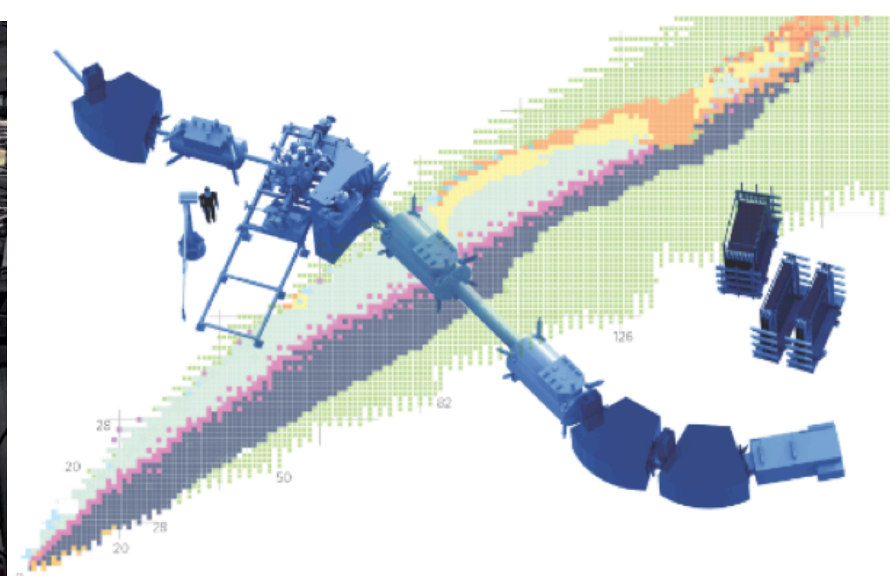

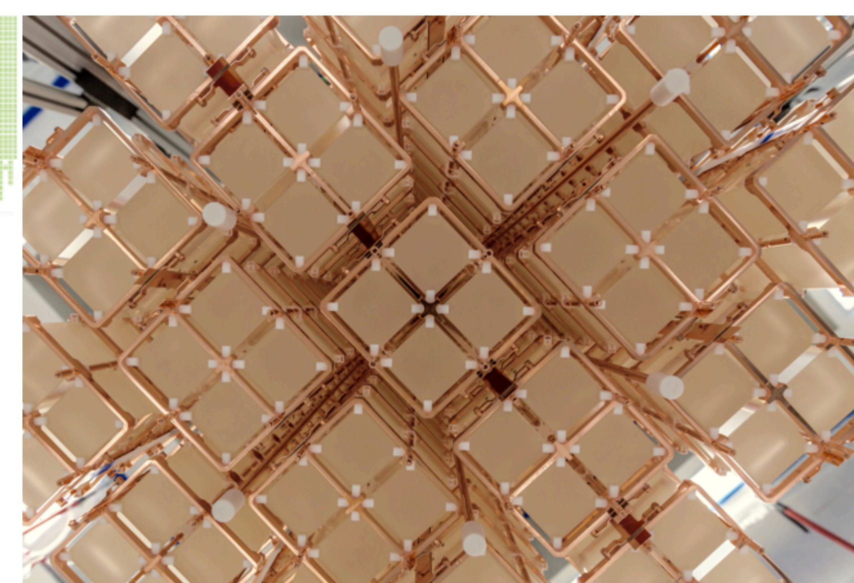

~ $1.2 Bn enterprise for construction+operations, ~4000 users

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Algorithms
From an Idea, through Development, to Production - Critical Element

Minimum time to solution ?

A.x=b

e.g. Multigrid

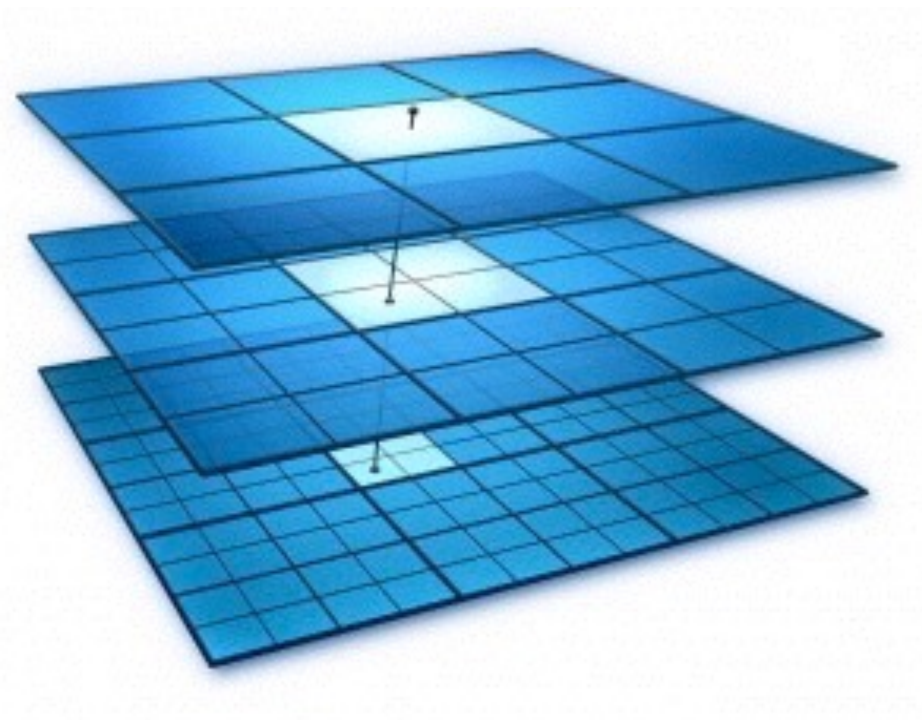

(USQCD, Joo and Edwards, JLab)

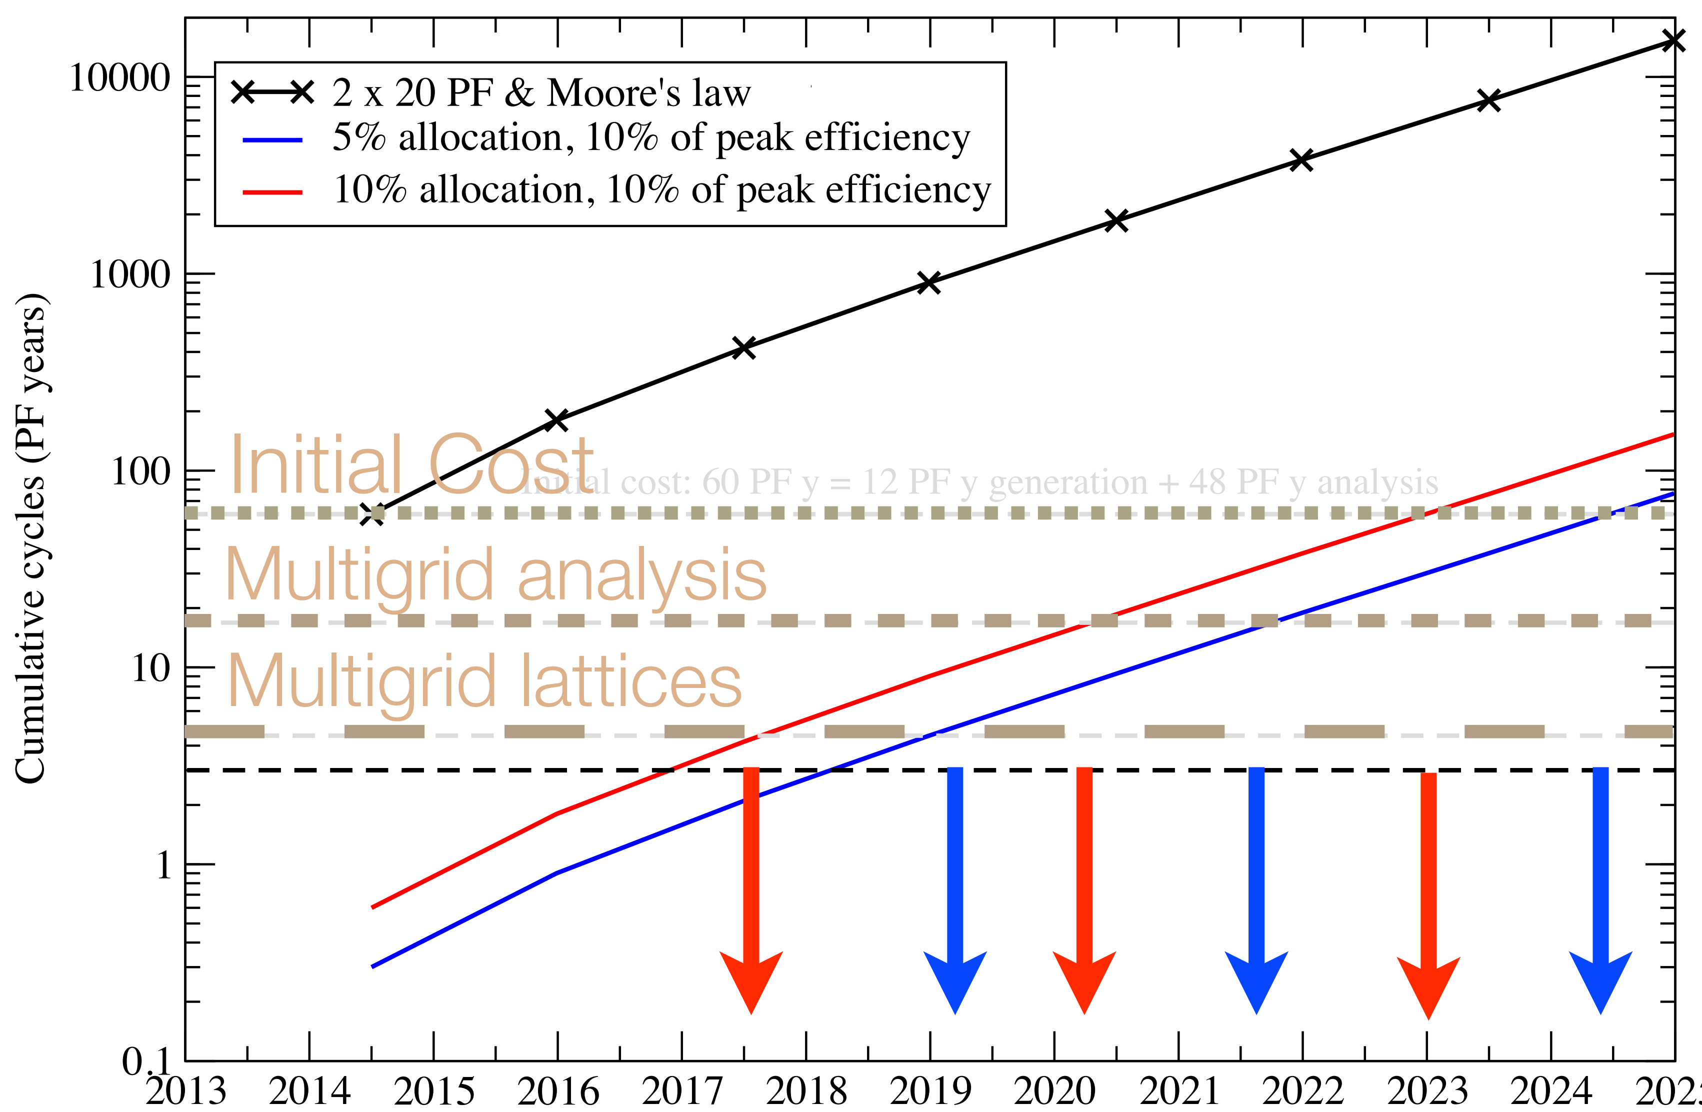

Essential, Extensive Collaboration

SciDAC Scientific Discovery through Advanced Computing

ECP EXASCALE COMPUTING PROJECT

USQCD-Nuclear Physics collaboration network

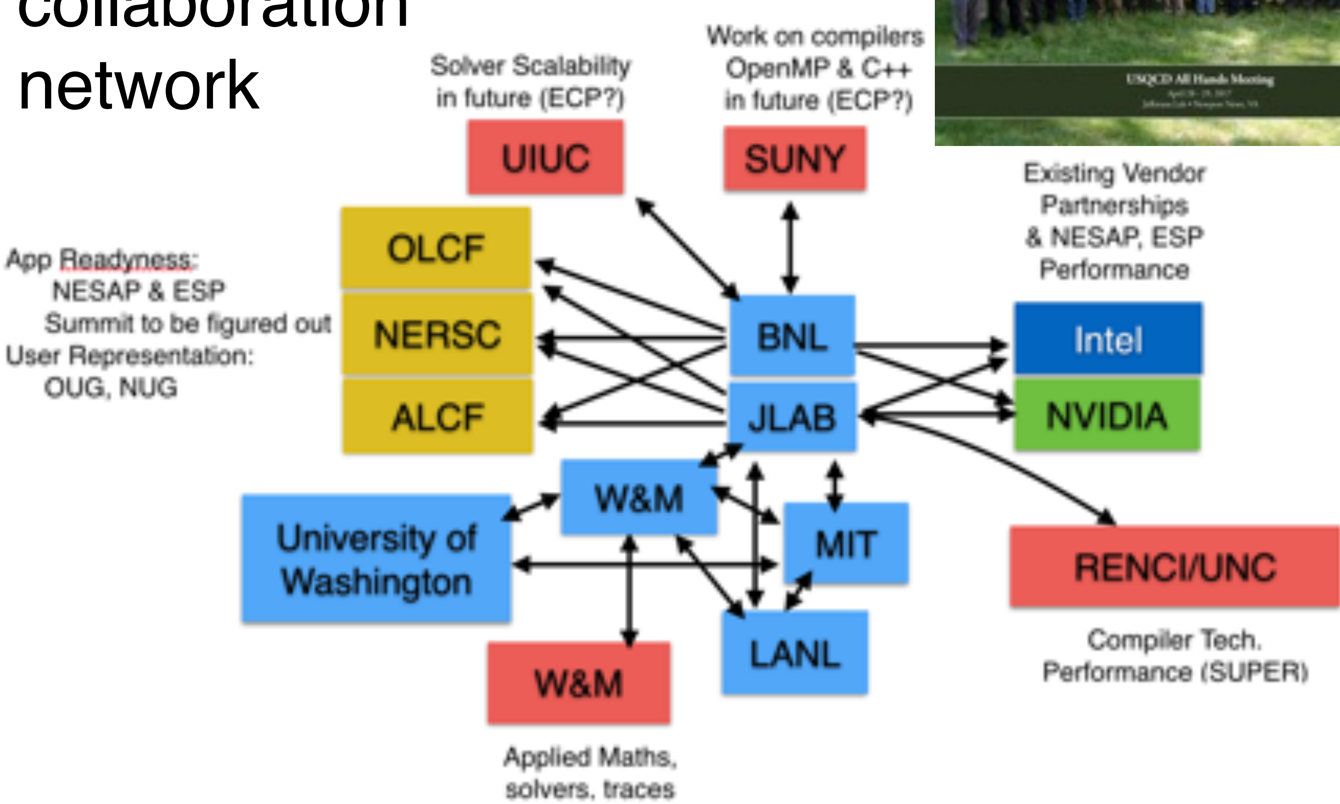

NUCLEI collaboration network

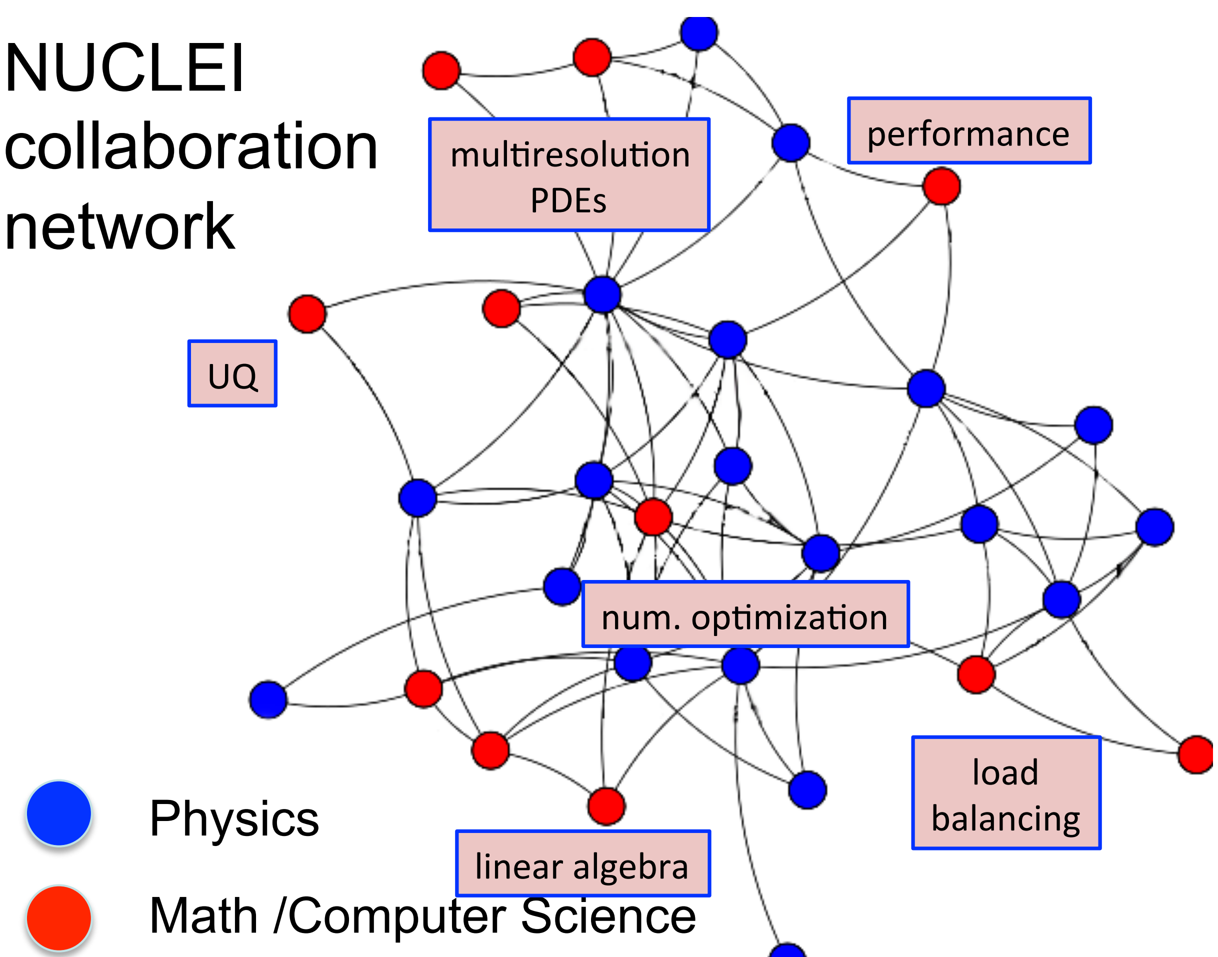

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Essential, Extensive Collaboration
Necessity is the Motherhood of Invention

2009

Mid-scale GPU-accelerated Cluster
Precursor to Exascale Architectures

QCDSP : 1998 Gordon Bell Prize

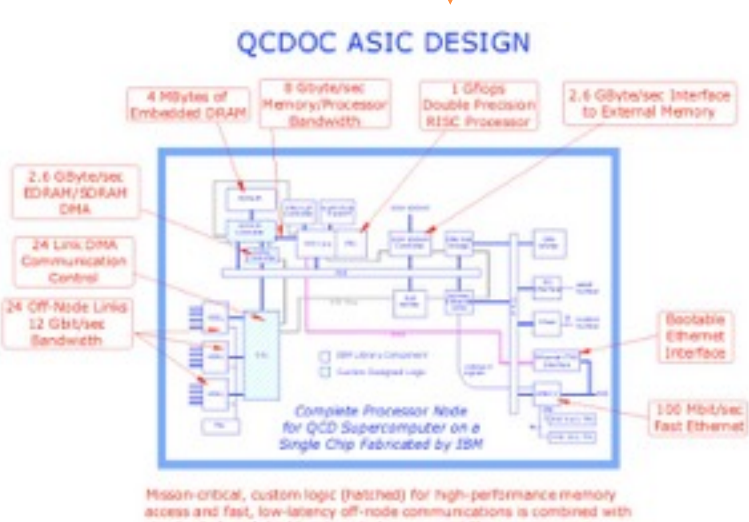

SciDAC
Scientific Discovery through Advanced Computing

Essential, Extensive Collaboration Porting Challenges

Code Portability is a currently a big issue, expect this to accelerate
Workflow is emerging as a major issue and also expected to accelerate

OLCF

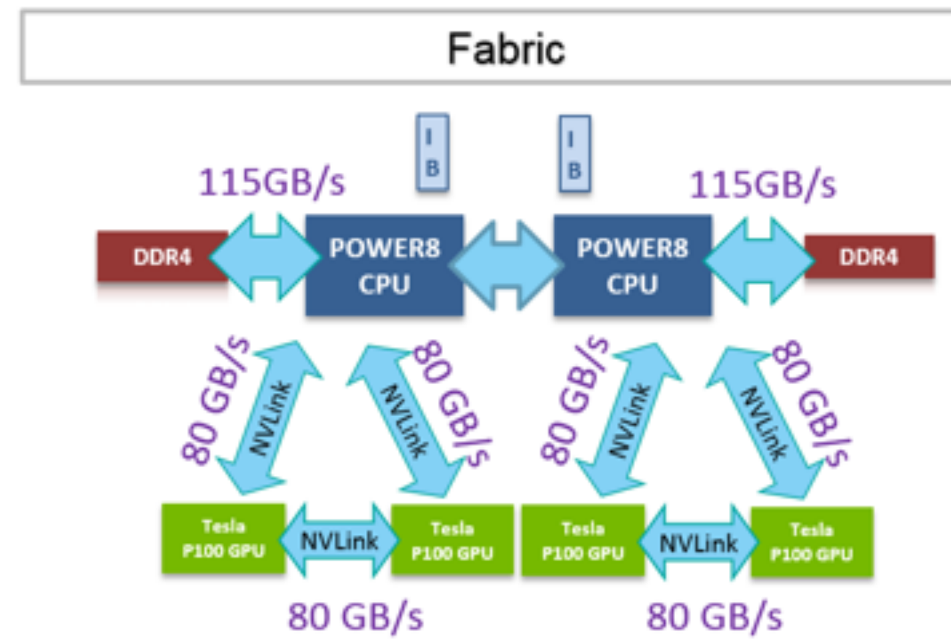

ALCF

Algorithms
From an Idea, through Development, to Production - Critical Element

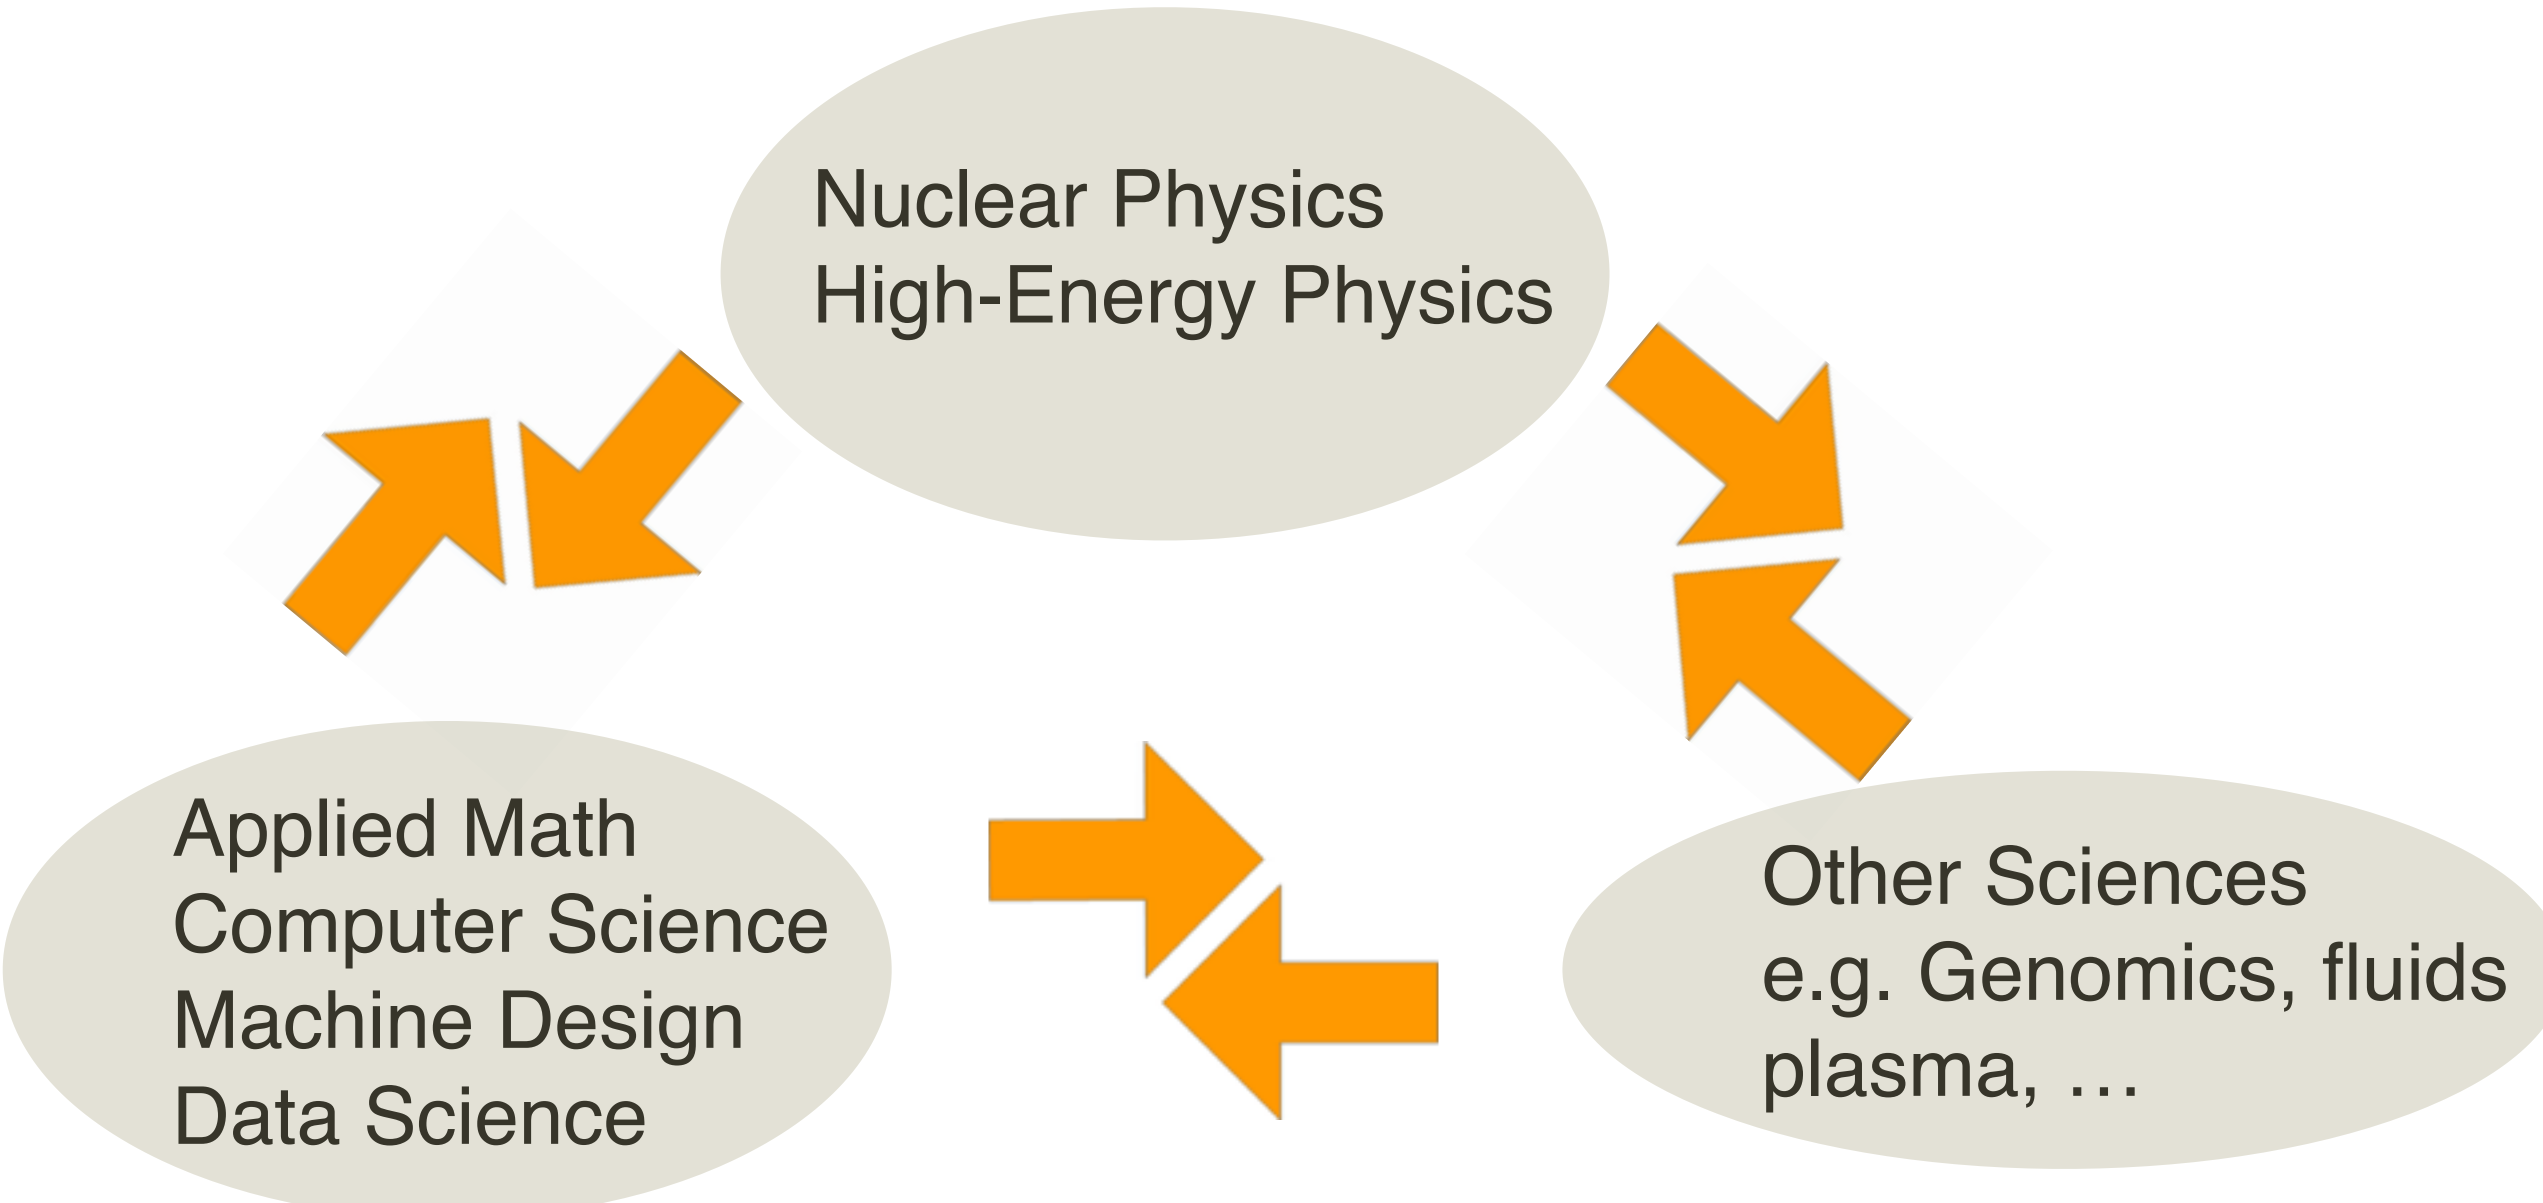

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Algorithms
From an Idea, through Development, to Production - Critical Element

e.g., MSU
-includes HEP, NP

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

(Commercial) Cloud Computing ?

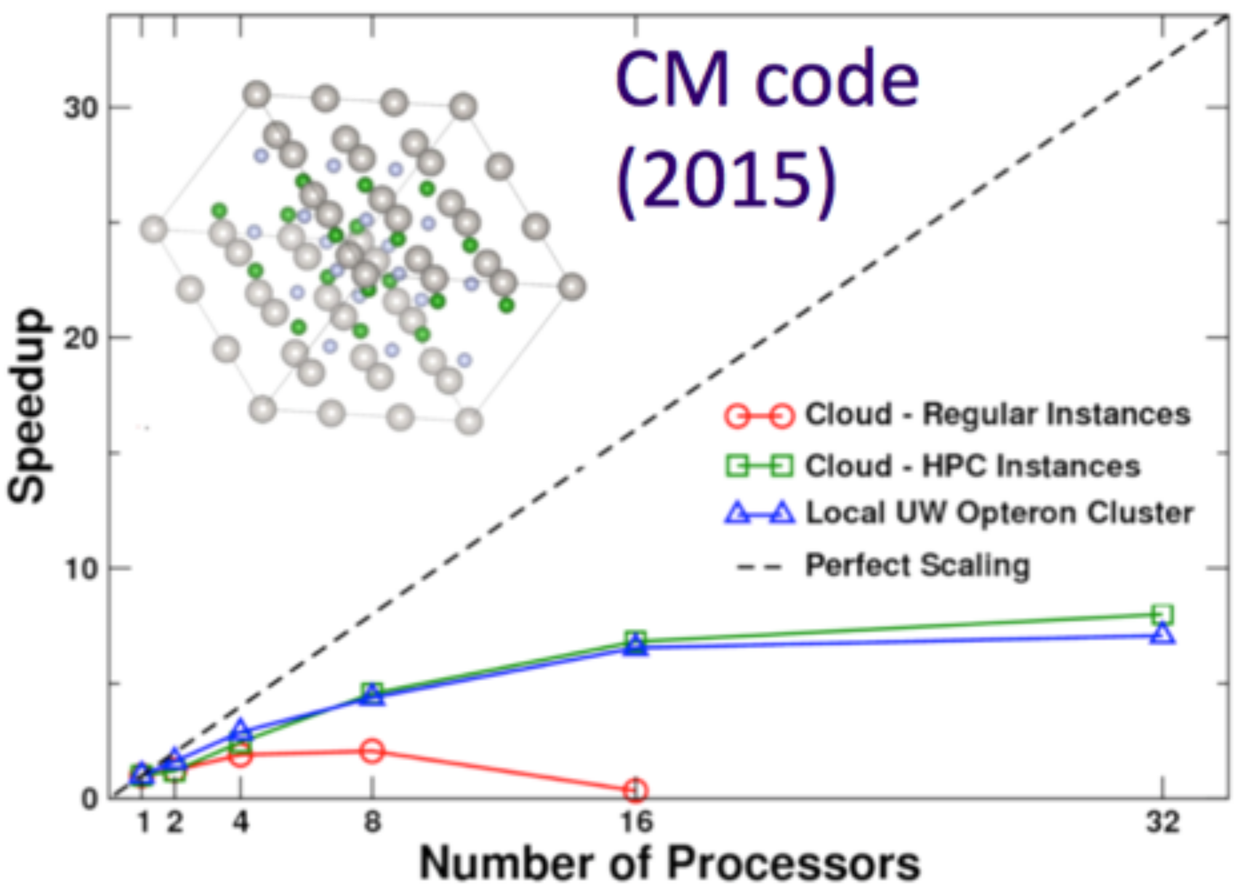

USQCD - ``designer'' code

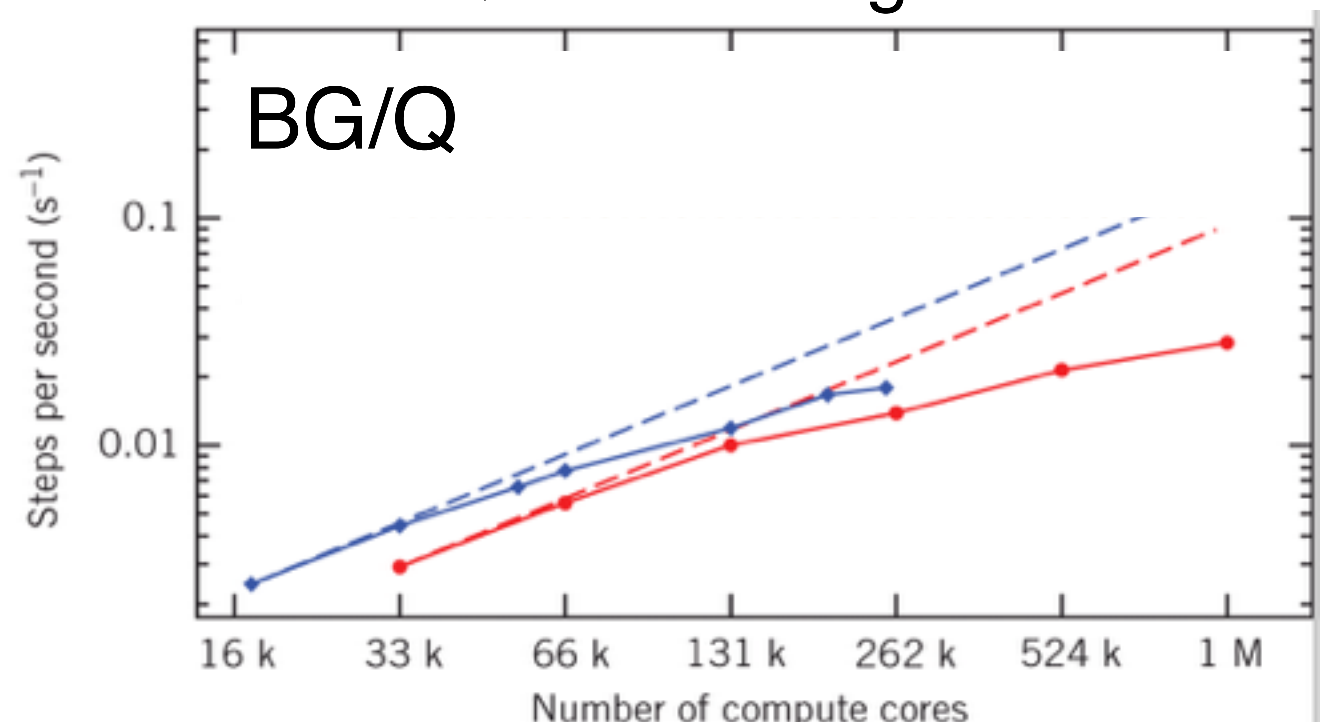

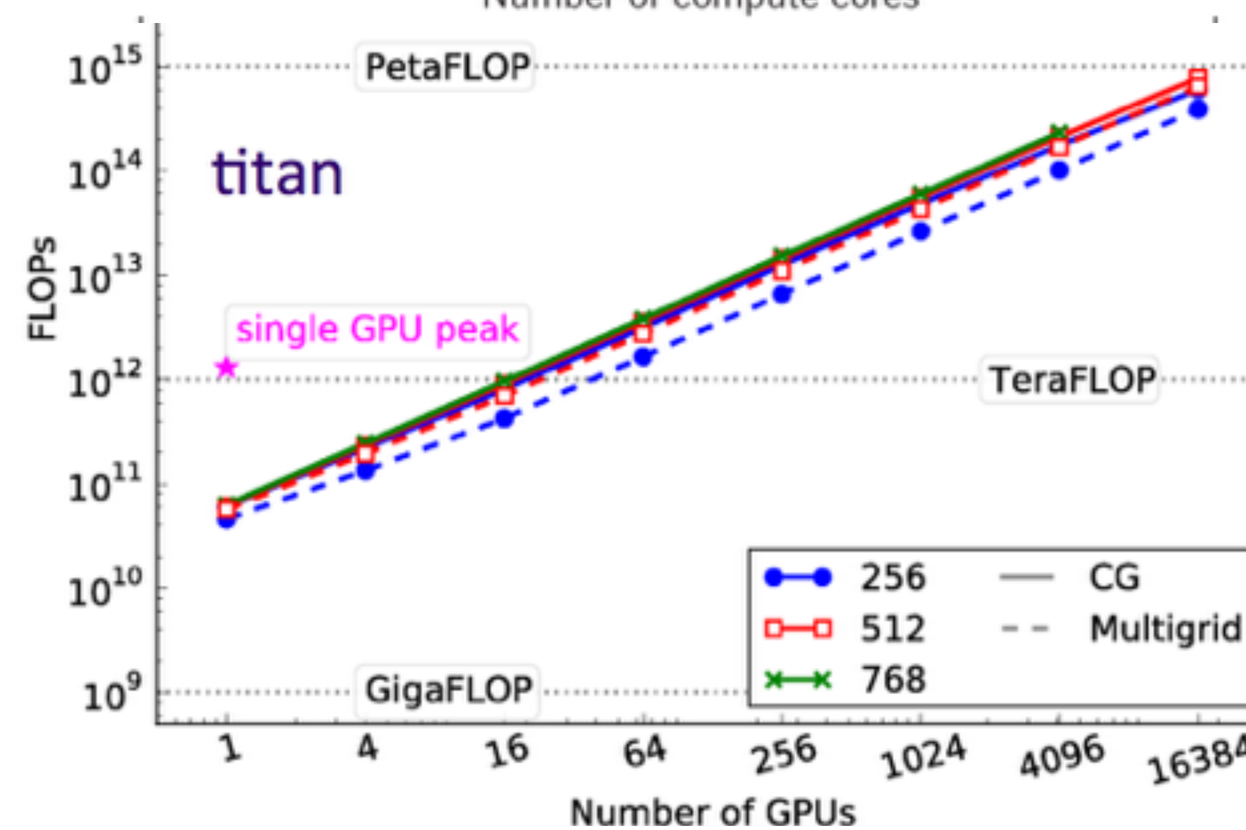

- Limited number of thread jobs are run in the cloud.
- Commercial cloud is > x5 expensive conpared to local dedicated clusters.

DOE-ASCR's Vision in the Exascale Era

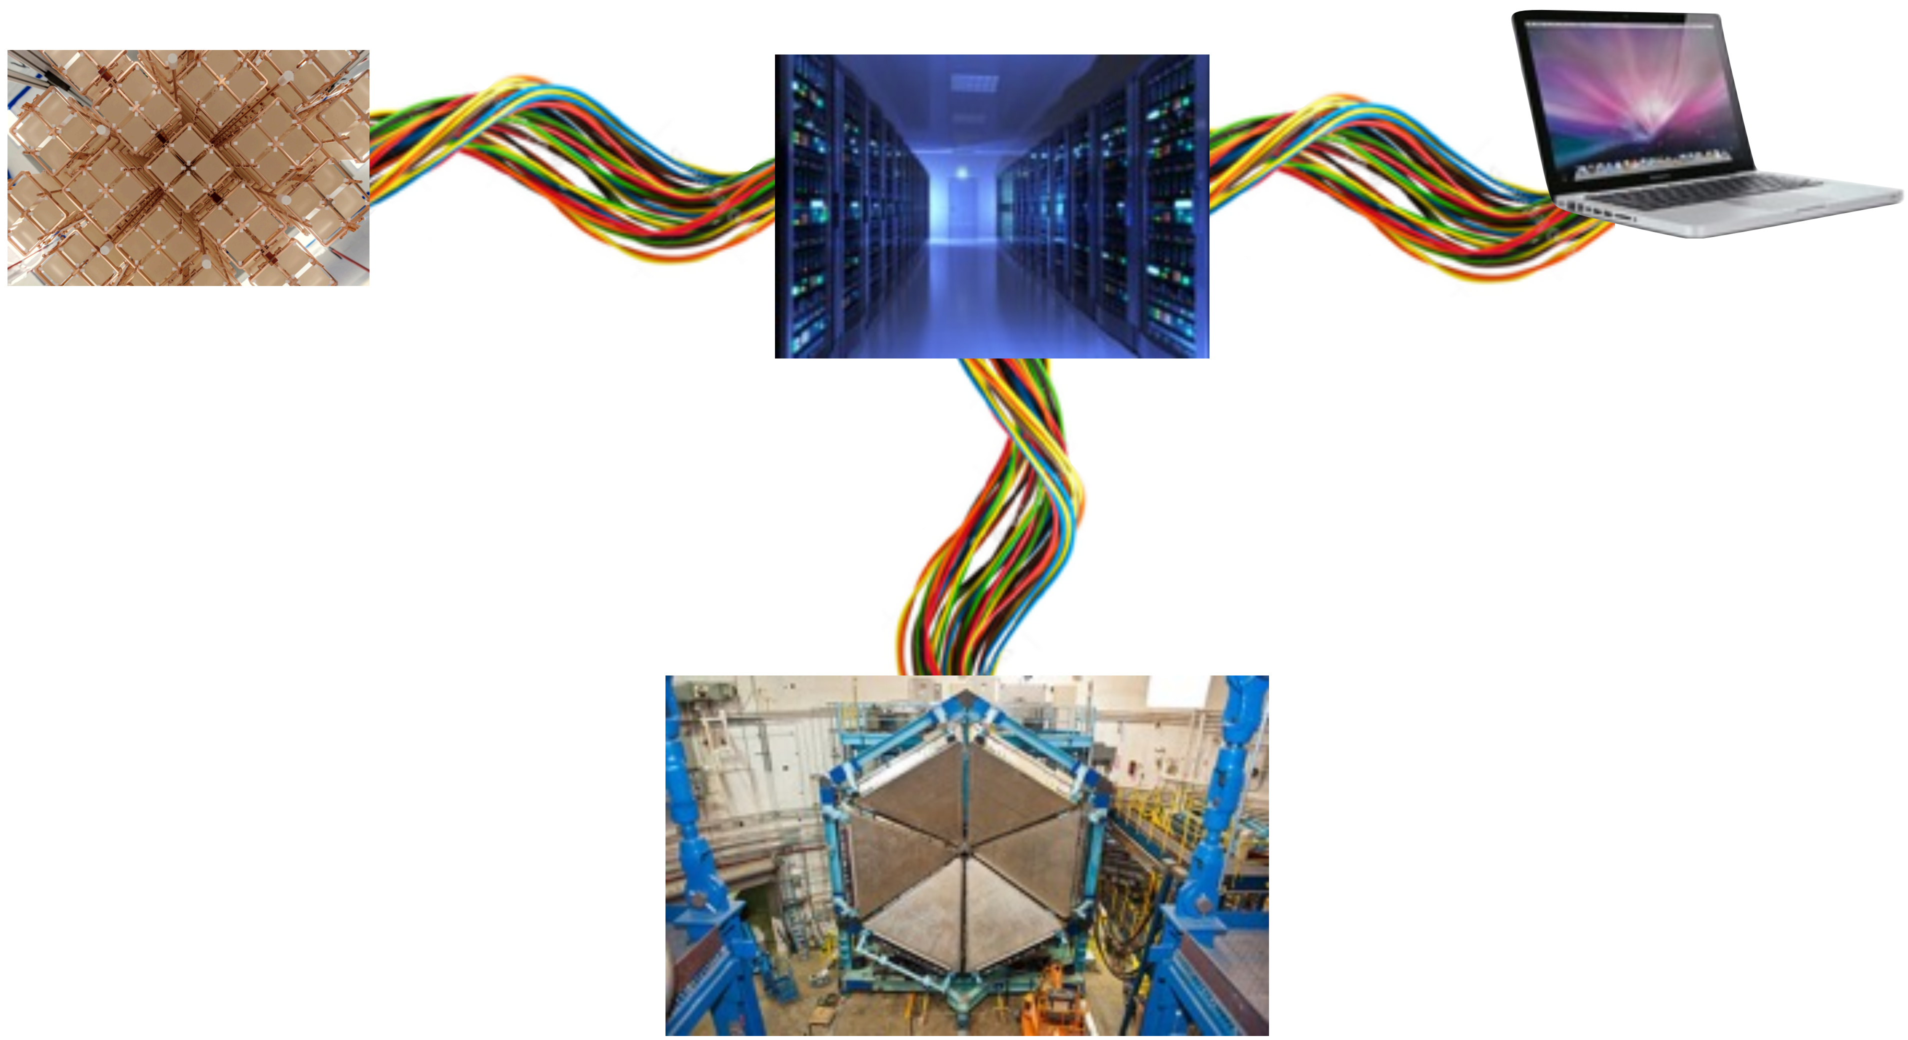

DOE-ASCR currently intends to provide Hardware:

- Small number of Exascale machines
- Big data pipes

Different from what NP currently has in place

- Efficacy for NP needs to be demonstrated

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

The NP community determined that during the Exascale Era and beyond, we will need to :

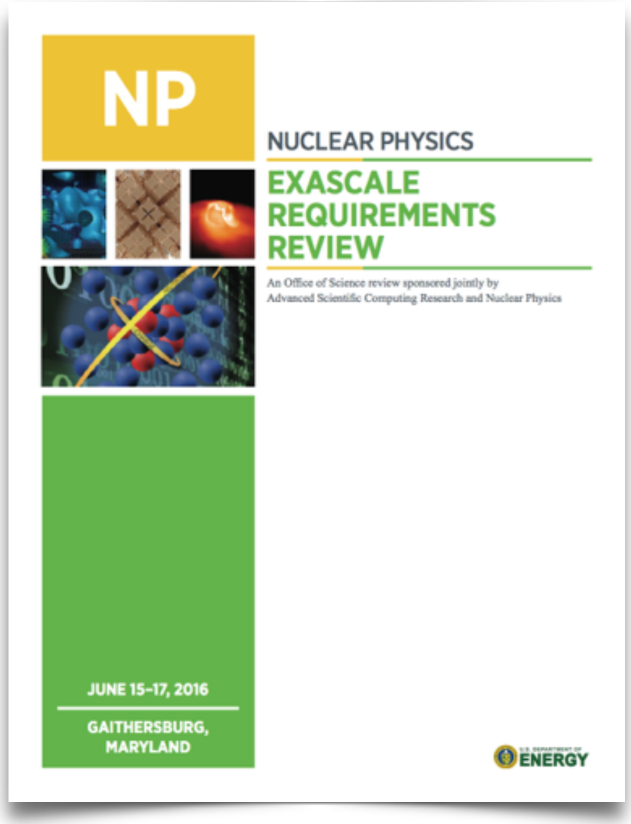

- Solve problems of extreme complexity and magnitude across a wide range of physical scales
- Explore parameter spaces and quantify uncertainties using ensembles of calculations
- Verify and validate algorithms, codes, and models
- Develop and optimize new codes, algorithms, models, and workflows
- Read, write, manage, analyze, curate, and track data of complexity and scale never before encountered
- Grow and sustain an appropriate workforce
  - enhanced collaboration between NP, ASCR, and NSF
  - creation of permanent positions at the interface between NP and ASCR
- Enhance local capacity hardware at laboratories and universities
  - integrating with ASCR infrastructure.

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025
NUCLEAR ASTROPHYSICS

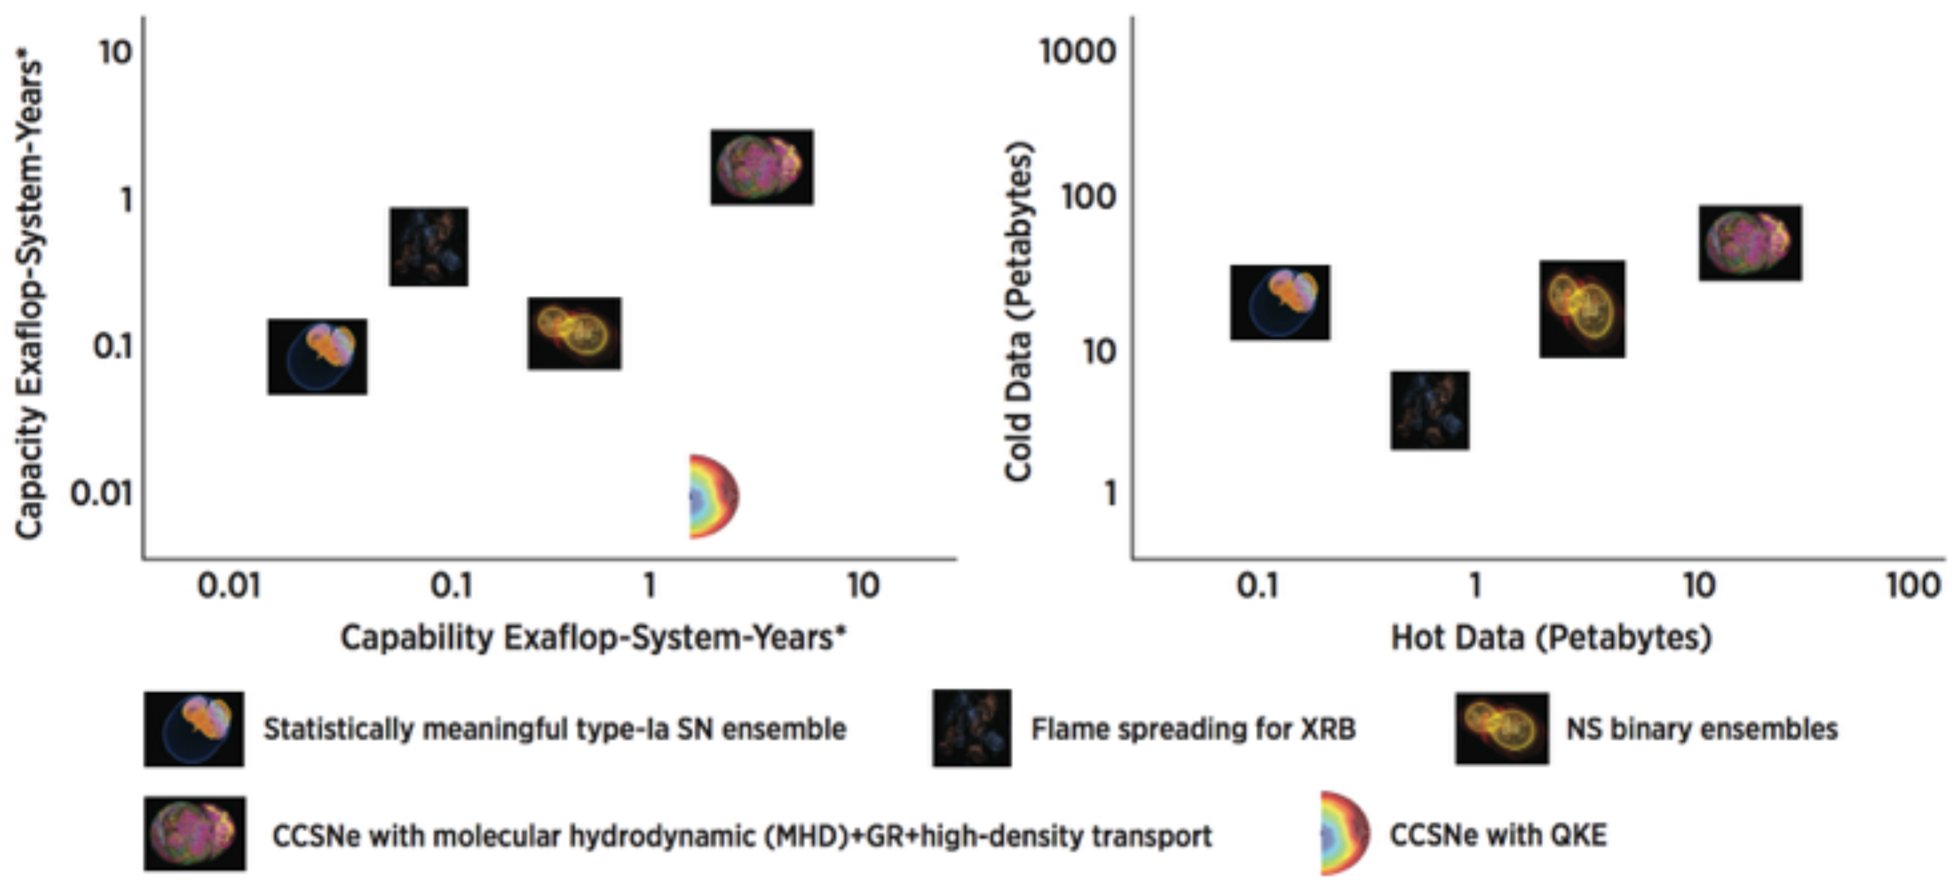

* Exaflop-system-year refers to the total amount of computation produced by an exascale computer in 1 year.

1 Exaflop-System is expected to sustain ~ 100 Petaflops on application code

Multi Exascale need in NP alone, >10x in the US

I expect that an Exascale machine will be immediately oversubscribed by a large factor

Exascale Ecosystem for Nuclear Physics

Solve computational problems of extreme complexity and magnitude across a very wide range of physical scales.

- Requires Exascale capability
  - hardware - with attributes consistent with NP code needs - diverse
    - large memory per node
    - memory bandwidth
    - low-latency communication
    - ….
  - software
  - applications
- Application readiness programs are in place : more=better
- New algorithms and mathematical formulations of our problems
  - new computing and runtime paradigms need exploration
- Build on and enhance SciDAC collaborations
- NP+HEP scientists need to remain fully engaged with hardware vendors
- Intra-site networks, HPC hardware, storage, HPC trained workforce, …

Exascale Ecosystem for Nuclear Physics

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Explore parameter spaces and quantify uncertainties using ensembles of calculations.

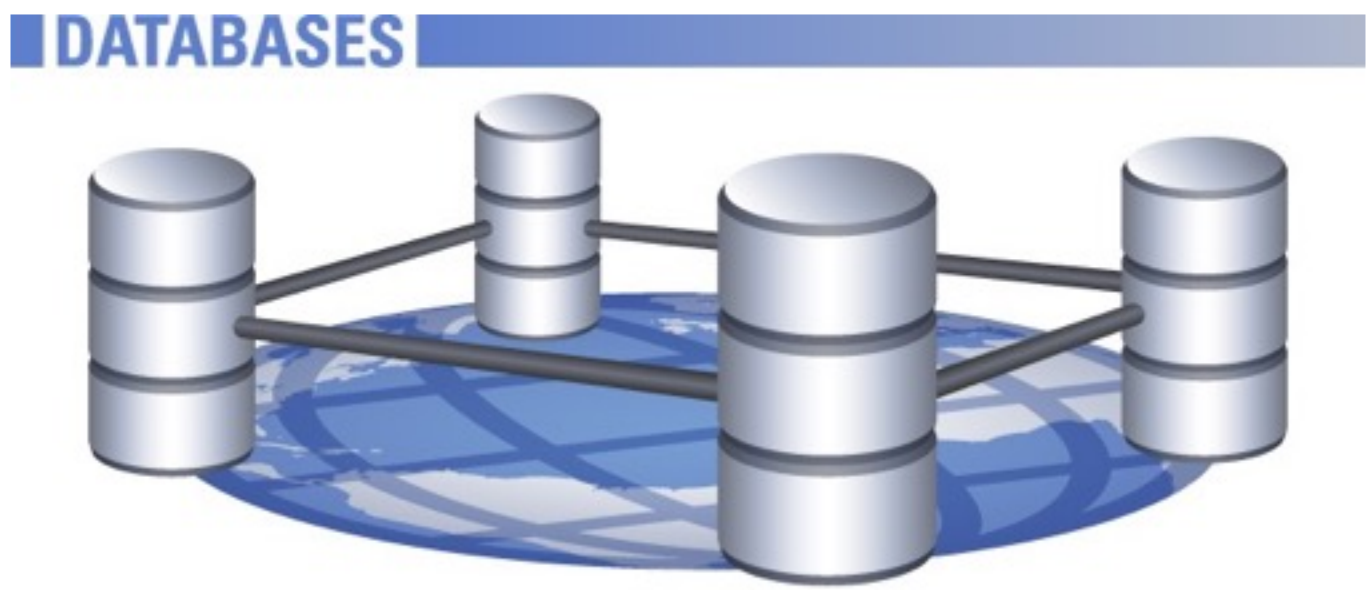

- Test theories and assumptions, find optimal solutions, propagate uncertainties, validate models, quantify uncertainties.
- Ensemble calculations and simulations, and the experimental program, require rapid-access and addressable databases, and spectrum of job sizes
- Databases required for data curation and long-term staging of hot and cold data for postproduction analysis — effort required in this area

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

Read, write, manage, analyze, curate, and track data/code of complexity and scale never before encountered

DOE: Data Management Plans

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

Read, write, manage, analyze, curate, and track data/code of complexity and scale never before encountered


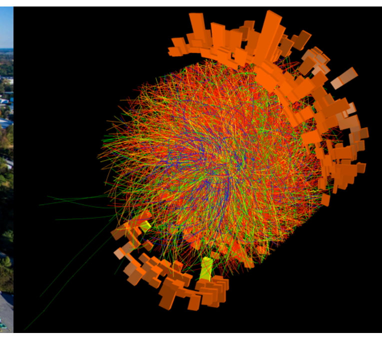
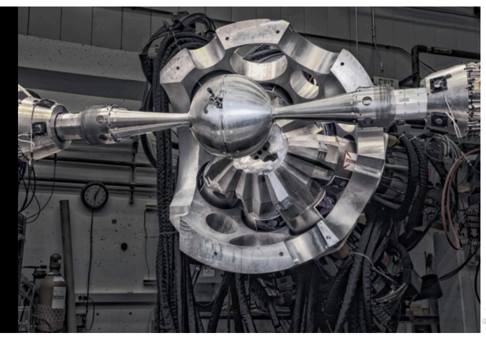
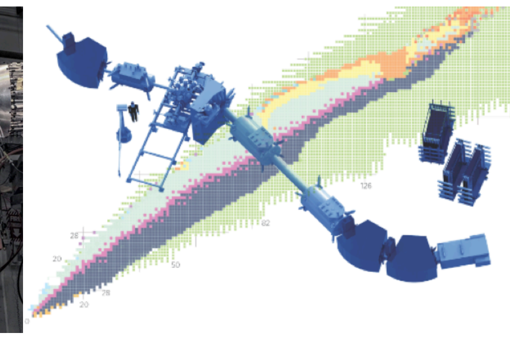
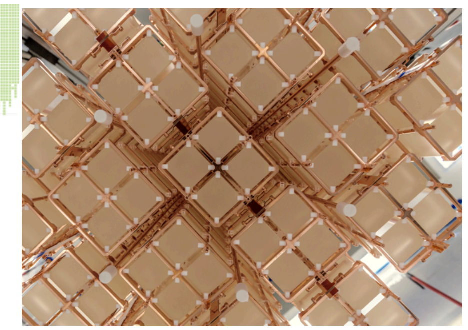

- Integration with existing and mature software, workflows, DAQs, movement and analysis
- Portability at all scales and locations
- Distinction between online and offline analysis is fading - real-time capabilities
- I/O becoming critical, collaborative efforts required
- Desirable for ASCR and NSF facilities to establish common attributes
  - authentication, software tools
  - implementation of data management plans and intellectual properties matters.
- Enhanced HPC expertise local to experiments/Labs required — porting to new architectures.

Exascale Ecosystem for Nuclear Physics

Grow and sustain a workforce to carry NP computational science through the exascale era, including enhanced collaboration between NP, ASCR, and NSF, and creation of permanent positions at the interface between NP and ASCR

All domains sciences recognize this as an issue in accomplishing their mission

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

Grow and sustain a workforce to carry NP computational science through the exascale era, including enhanced collaboration between NP, ASCR, and NSF, and creation of permanent positions at the interface between NP and ASCR

- SciDAC : PhD students, PostDocs, junior senior scientists
  - develop code and algorithms on evolving capability and capacity hardware.
    - without SciDAC, codes would not be running well on GPU, KNL or BG/LPQ
  - SciDAC, NESAP, CAAR, ESP and ECP in place
  - Other training at the facilities and universities

- Growing and sustaining a HPCD workforce through exascale is vital to NP mission
  - workforce development should include enhanced collaboration between NP, ASCR, and NSF and new positions at the ASCR/NP interface.

- University and Laboratory positions are essential

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Exascale Ecosystem for Nuclear Physics

Enhance local capacity hardware at laboratories and universities that is capable of efficiently integrating with ASCR infrastructure and facilities.

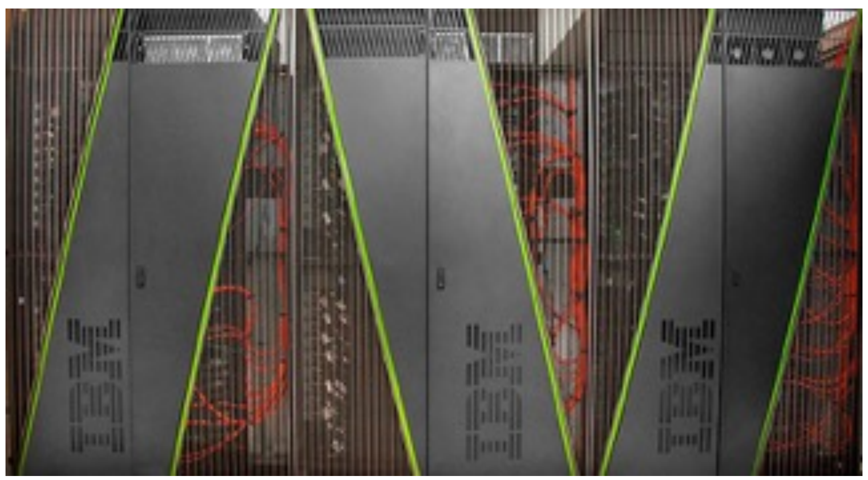

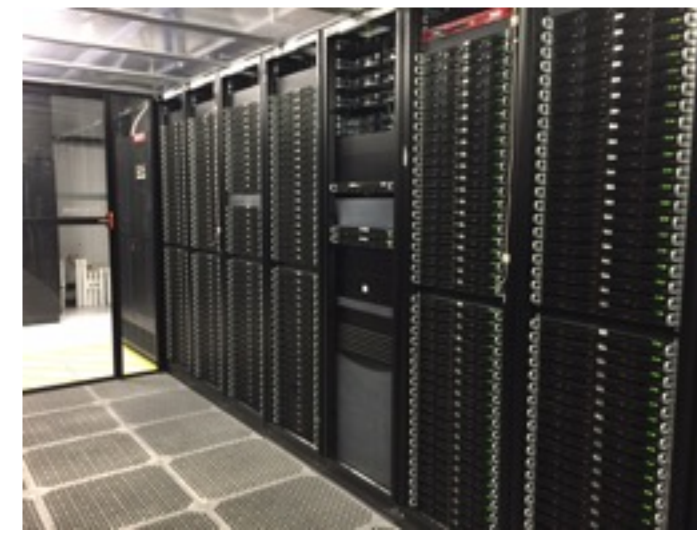

- Critical element of Exascale Ecosystem
- 50% of USQCD compute resources today
  - Leadership Class facilities oversubscribed by large factors
- Hardware optimized for scientific task
- Code and Algorithm development
- Speed of Science
  - I can run 1 job exploring a new idea on 600 cores for 3 weeks continually immediately on local cluster if I want
- Training tool, student and Postdoctoral
- ``Captures" students
- Historically has led to HPC evolution

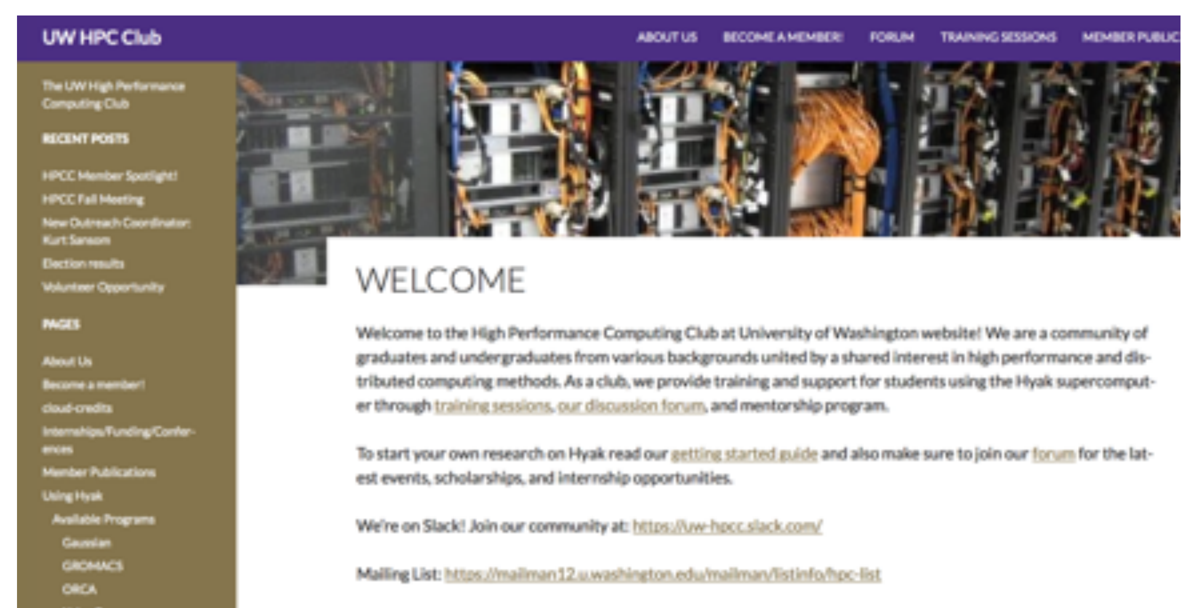

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Top 500 List

	NAME	SPECS	SITE	COUNTRY	CORES	RMAX PFLOP/S	POWER MW
1	Sunway TaihuLight	Shenwei SW26010 (260C 1.45 GHz) Custom interconnect	NSCC in Wuxi	China	10,649,600	93.0	15.4
2	Tianhe-2 (Milkyway-2)	Intel Ivy Bridge (12C 2.2 GHz) & Xeon Phi (57C 1.1 GHz), Custom interconnect	NSCC in Guangzhou	China	3,120,000	33.9	17.8
3	Titan	Cray XK7, Opteron 6274 (16C 2.2 GHz) + Nvidia Kepler GPU, Custom interconnect	DOE/SC/ORNL	USA	560,640	17.6	8.2
4	Sequoia	IBM BlueGene/Q, Power BQC (16C 1.60 GHz), Custom interconnect	DOE/NNSA/LLNL	USA	1,572,864	17.2	7.9
5	Cori	Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect	DOE/SC/LBNL	USA	622,336	14.0	

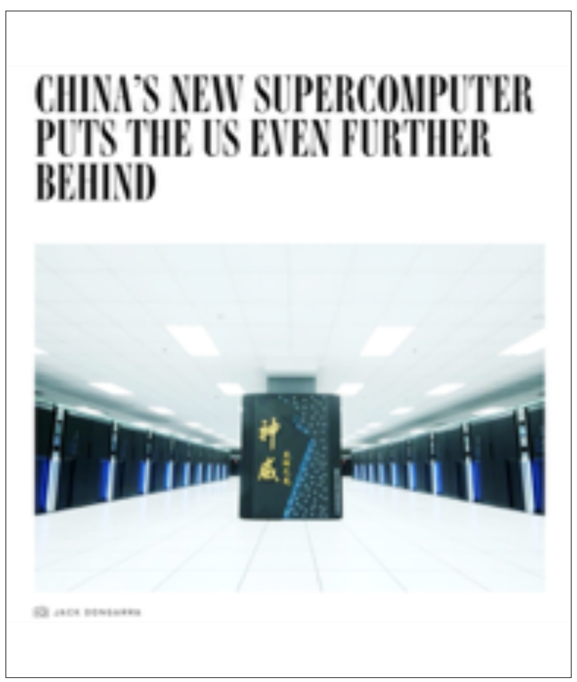
CHINA'S NEW SUPERCOMPUTER PUTS THE US EVEN FURTHER BEHIND

Made in China !

~125 Petaflops, ~15 MWatts

~11 million compute cores, ~1.3 PB ram

US will deploy ~100 Petaflop machines ~ 2017-18

expect ~ 1 Exaflop machines ~ 2022

Top 500 List

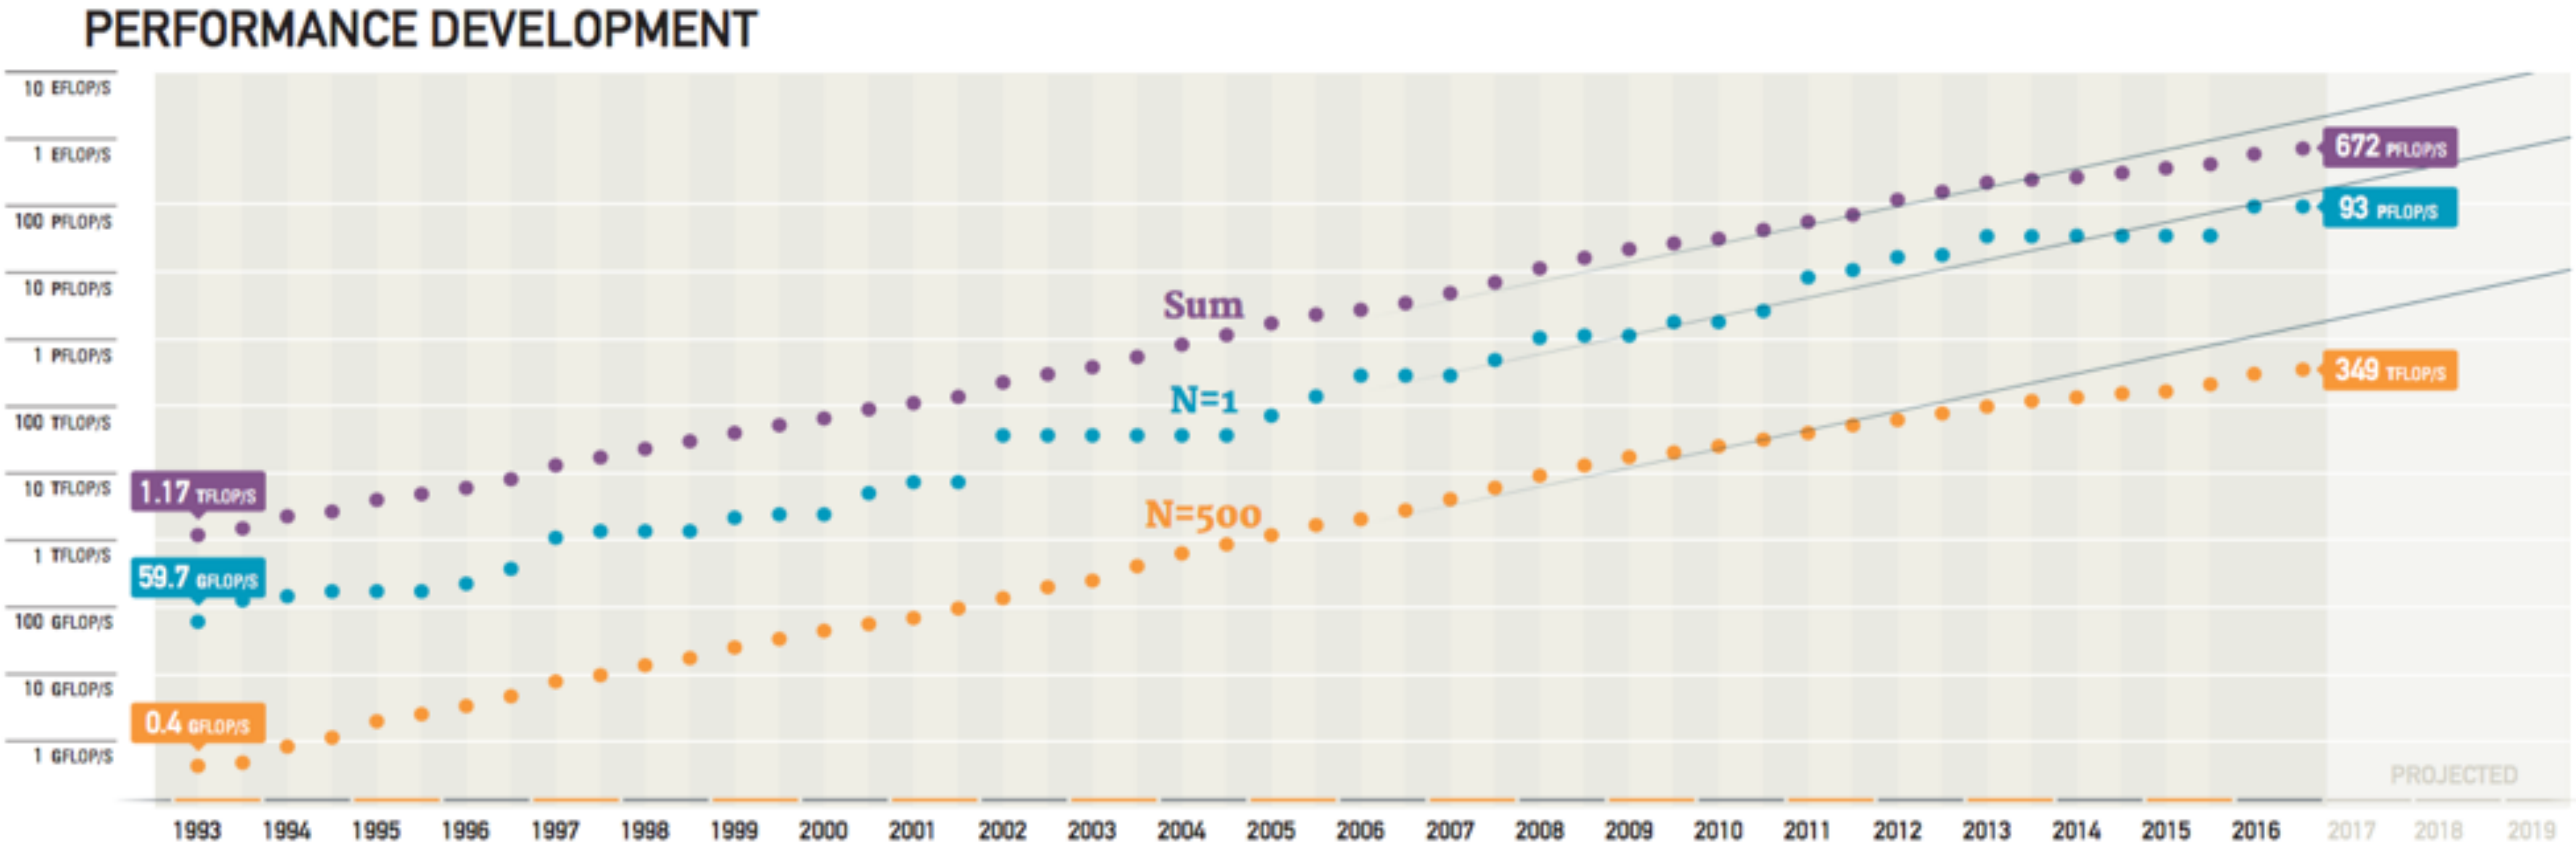
PERFORMANCE DEVELOPMENT
10 EFLOP/S
1 EFLOP/S
100 PFLOP/S
10 PFLOP/S
1 PFLOP/S
100 TFLOP/S
10 TFLOP/S
1 TFLOP/S
100 GFLOP/S
10 GFLOP/S
1 GFLOP/S
Sum
N=1
N=500
1.17 TFLOP/S
59.7 GFLOP/S
0.4 GFLOP/S
672 PFLOP/S
93 PFLOP/S
349 TFLOP/S
PROJECTED
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019

Expectations

2009

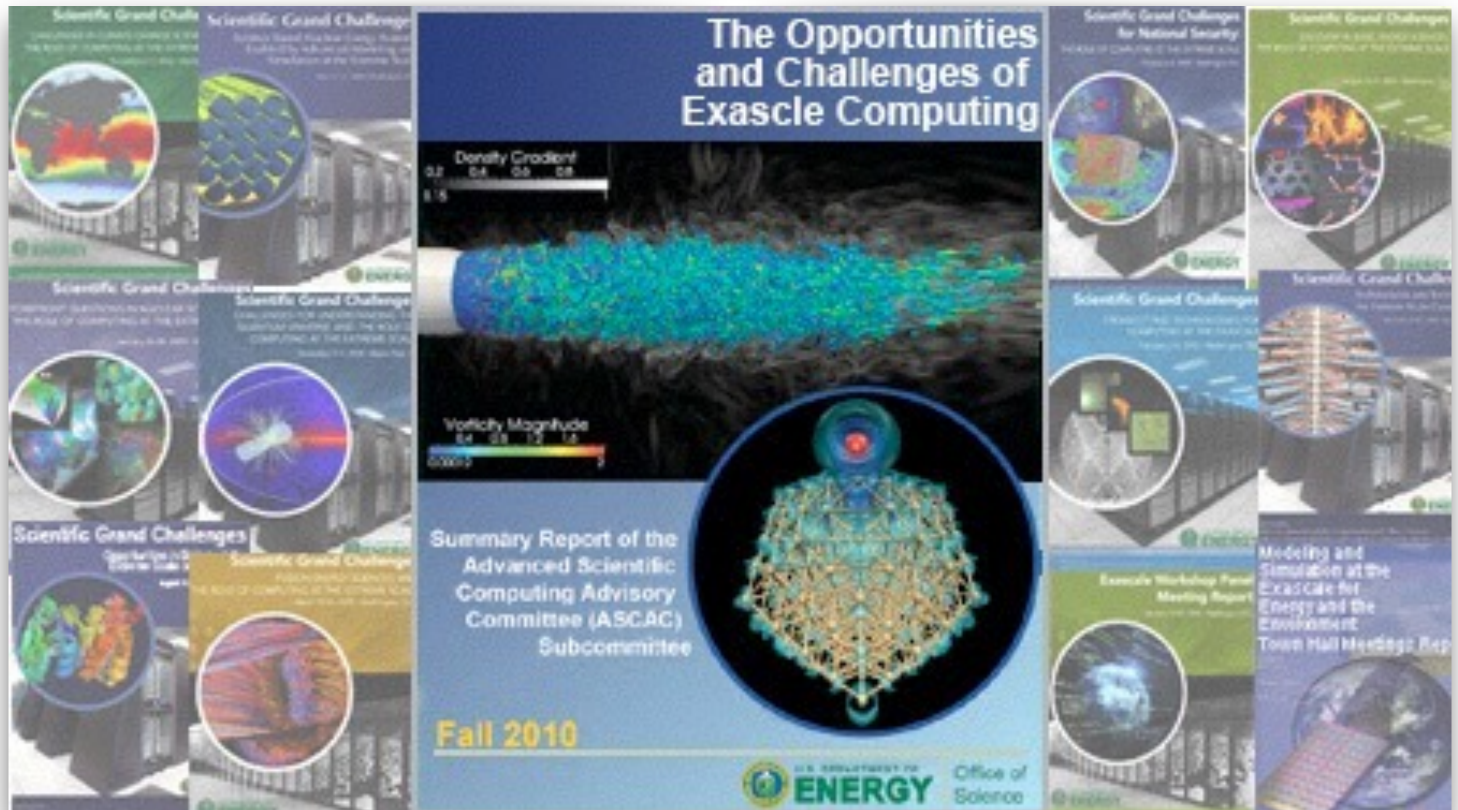

2016

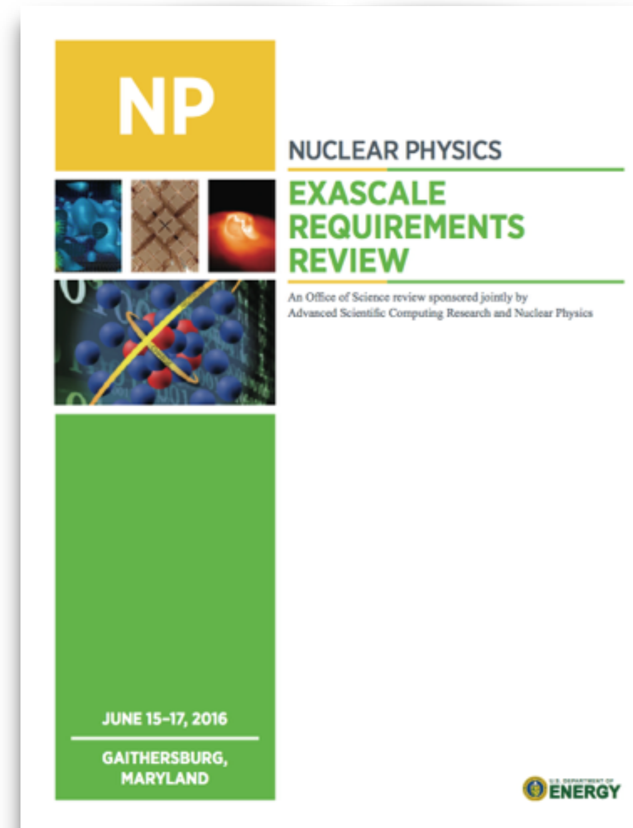

Reality:

- Years to Exascale in 2010 = 7, in 2017 = 4
  - increasing demand by other fields, within the field
  - competing well for resources
    - evolving landscape with non-US competition for US compute resources

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

New (disruptive?) Technologies

Quantum Computing

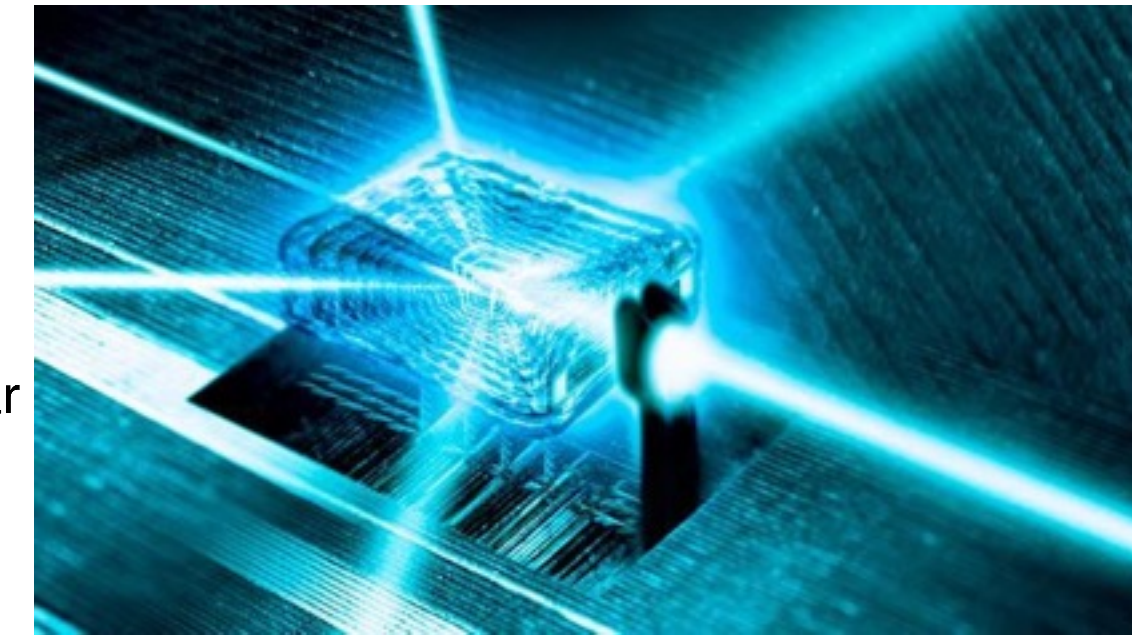

- NP has some computations that require exponential time on a classical computer - quantum computer?
- Microsoft, Google, IBM indicate ~50 qubit QC's up within a year
- NP could start exploring ….

Quantum Testbeds Stakeholder Workshop

Mayflower Hotel
1127 Connecticut Avenue NW
Washington, DC 20036
February 14 – 16, 2017

U.S. DEPARTMENT OF ENERGY

Machine Learning

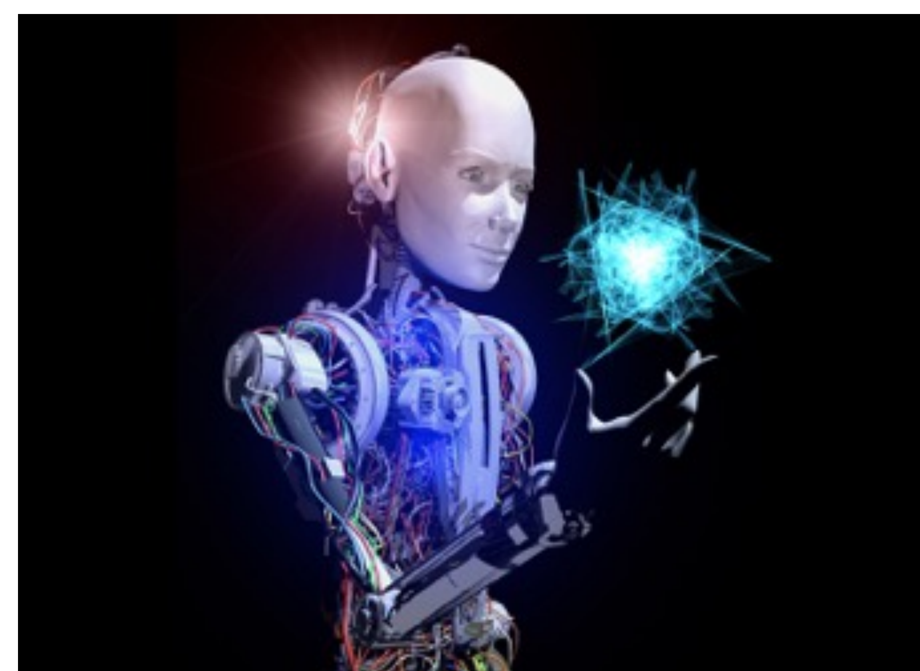

- Already in use in HEP expt at some level
- Sophisticated pattern recognition - correlations in data
- Nonlinear regression on steroids+epo
- Some NP engagement, and much more, would likely be valuable

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

Summary

The next decade will be looked back upon as a truly astonishing period in nuclear physics and in our understanding of fundamental aspects of nature.

This will be made possible by advances in scientific computing and in how the NP community organizes and collaborates, and how DOE and NSF supports this, to take full advantage of these advances.