Strategy for Computational NP in the LRP: Some Possibilities

- Science case: coupling to facilities and to the high priority projects that the Town Meetings identify

- Mapping onto DOE priorities
  - DOE NP/ASCR partnerships
  - SciDAC
  - Laboratory partnerships and the 2017 machines

- What to request: what might be plausibly accomplished
  - funding as a "third field" - the research budget problem
  - joint initiatives with partners
  - workforce development

An Opportunity for Nuclear Physics

- Established NP computational programs impacting all aspects of NP experiment
- Computational NP research aligns with a DOE and national high priority
- The NP community has demonstrated an ability to organize and to build the kinds of partnerships with ASCR necessary for success
- Co-location: strong NP groups already exist at the laboratories who host the petascale machines and the supporting applied math/CS groups

Argonne	Mira	Quantum Monte Carlo
Berkeley	NERSC/Edison	Nuclear Astro, HI, Lattice QCD
Livermore	Sequoia	Lattice QCD, Nuclear Structure
Los Alamos/Sandia	Celio	Quantum Monte Carlo, Nuclear Astro
Oak Ridge	Titan	Nuclear Structure, Nuclear Astro

- Strong facility-relevant computational programs established at JLab, BNL

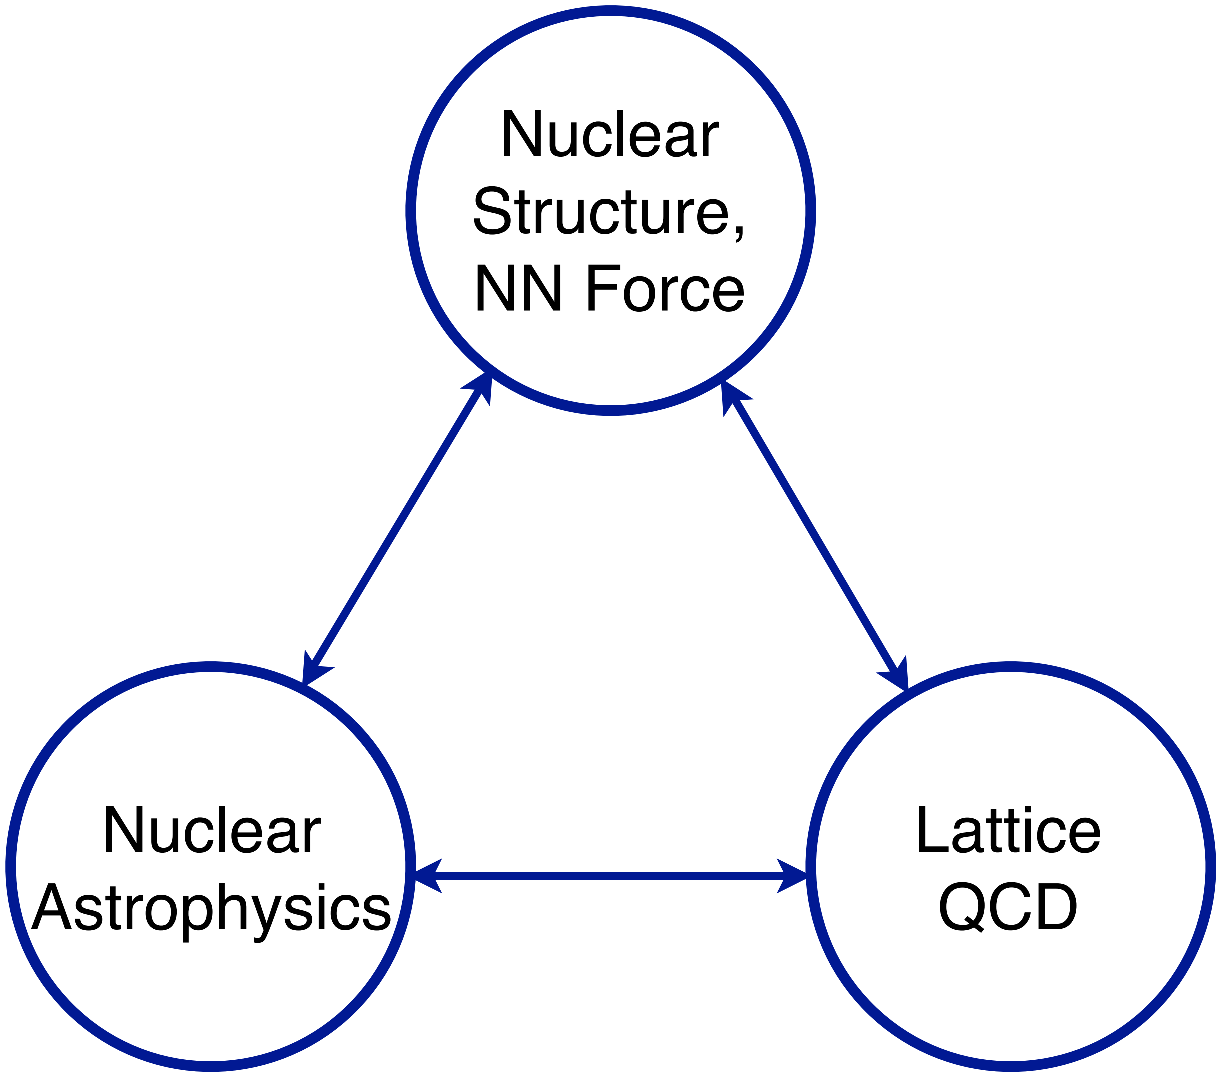
Nuclear
Structure,
NN Force
Nuclear
Astrophysics
Lattice
QCD

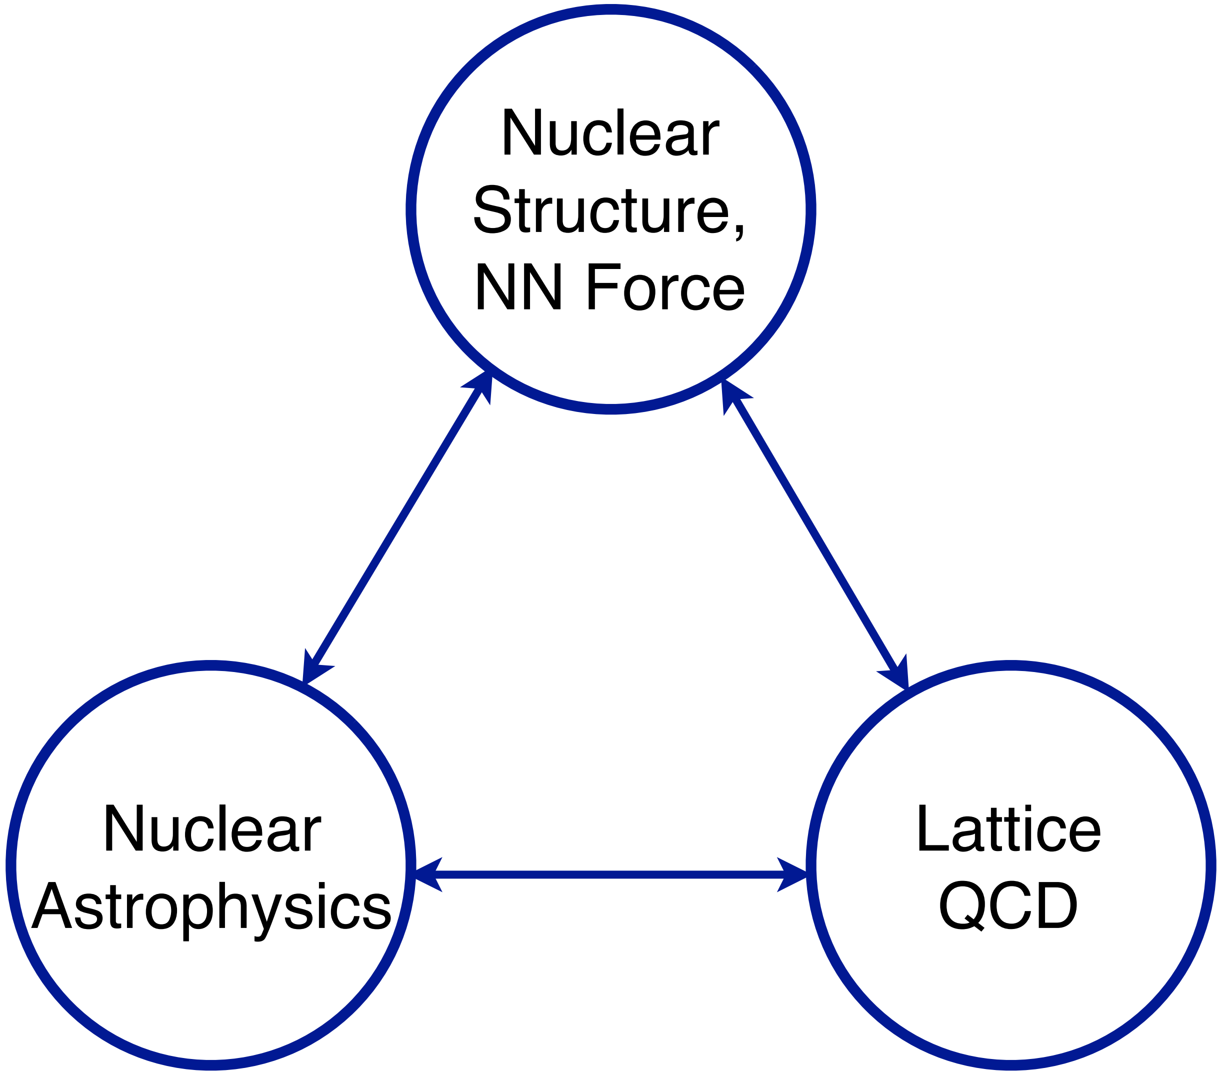

RHIC: QCD plasma

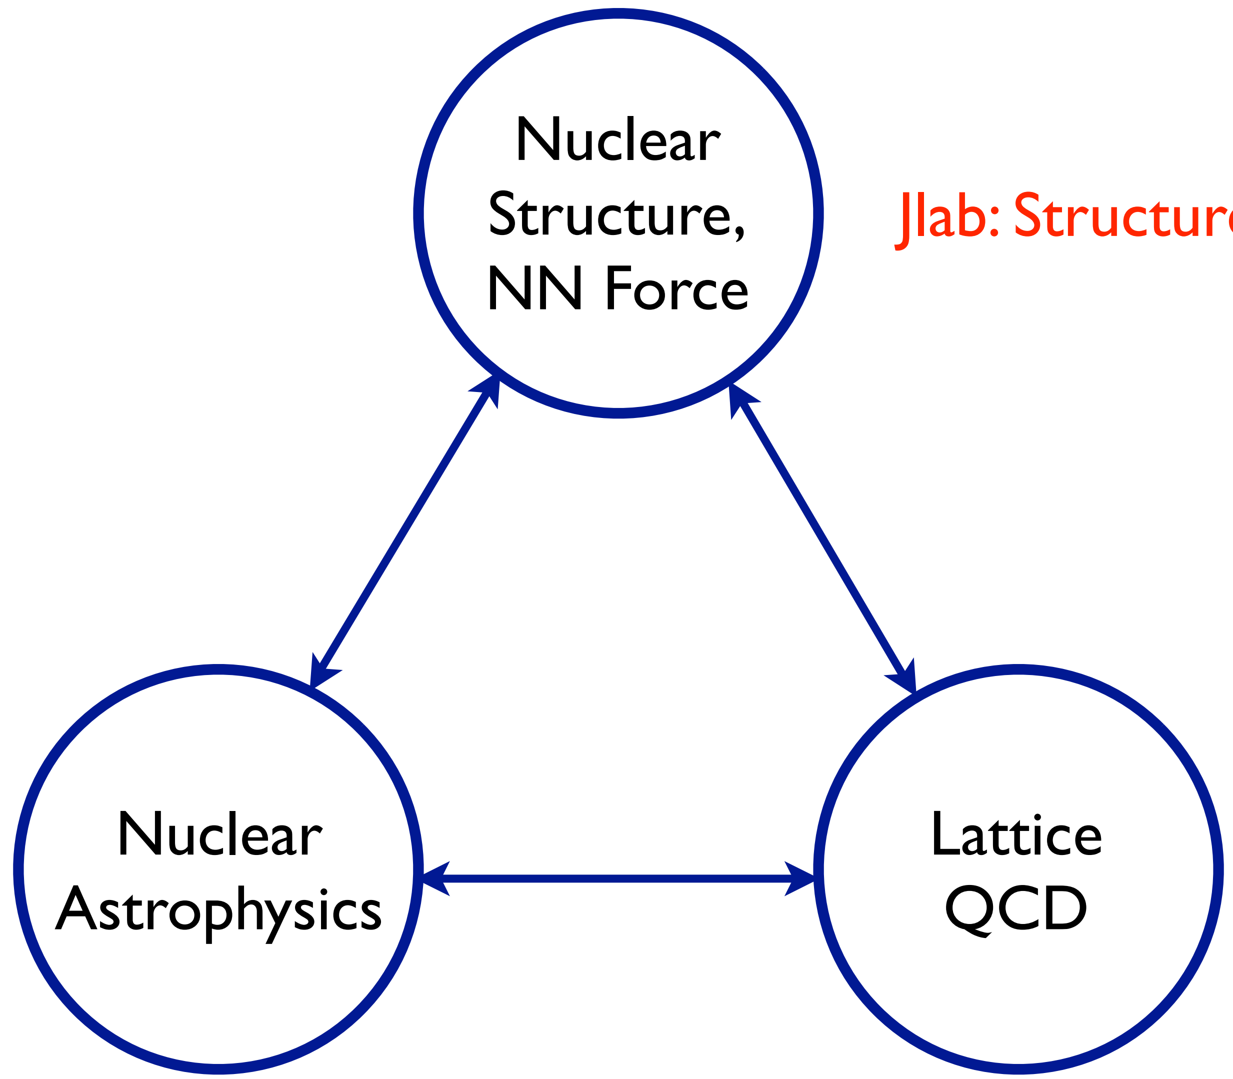

Jlab: Structure/Currents

RHIC: QCD plasma

Jlab/RHIC spin: Nucleon Structure

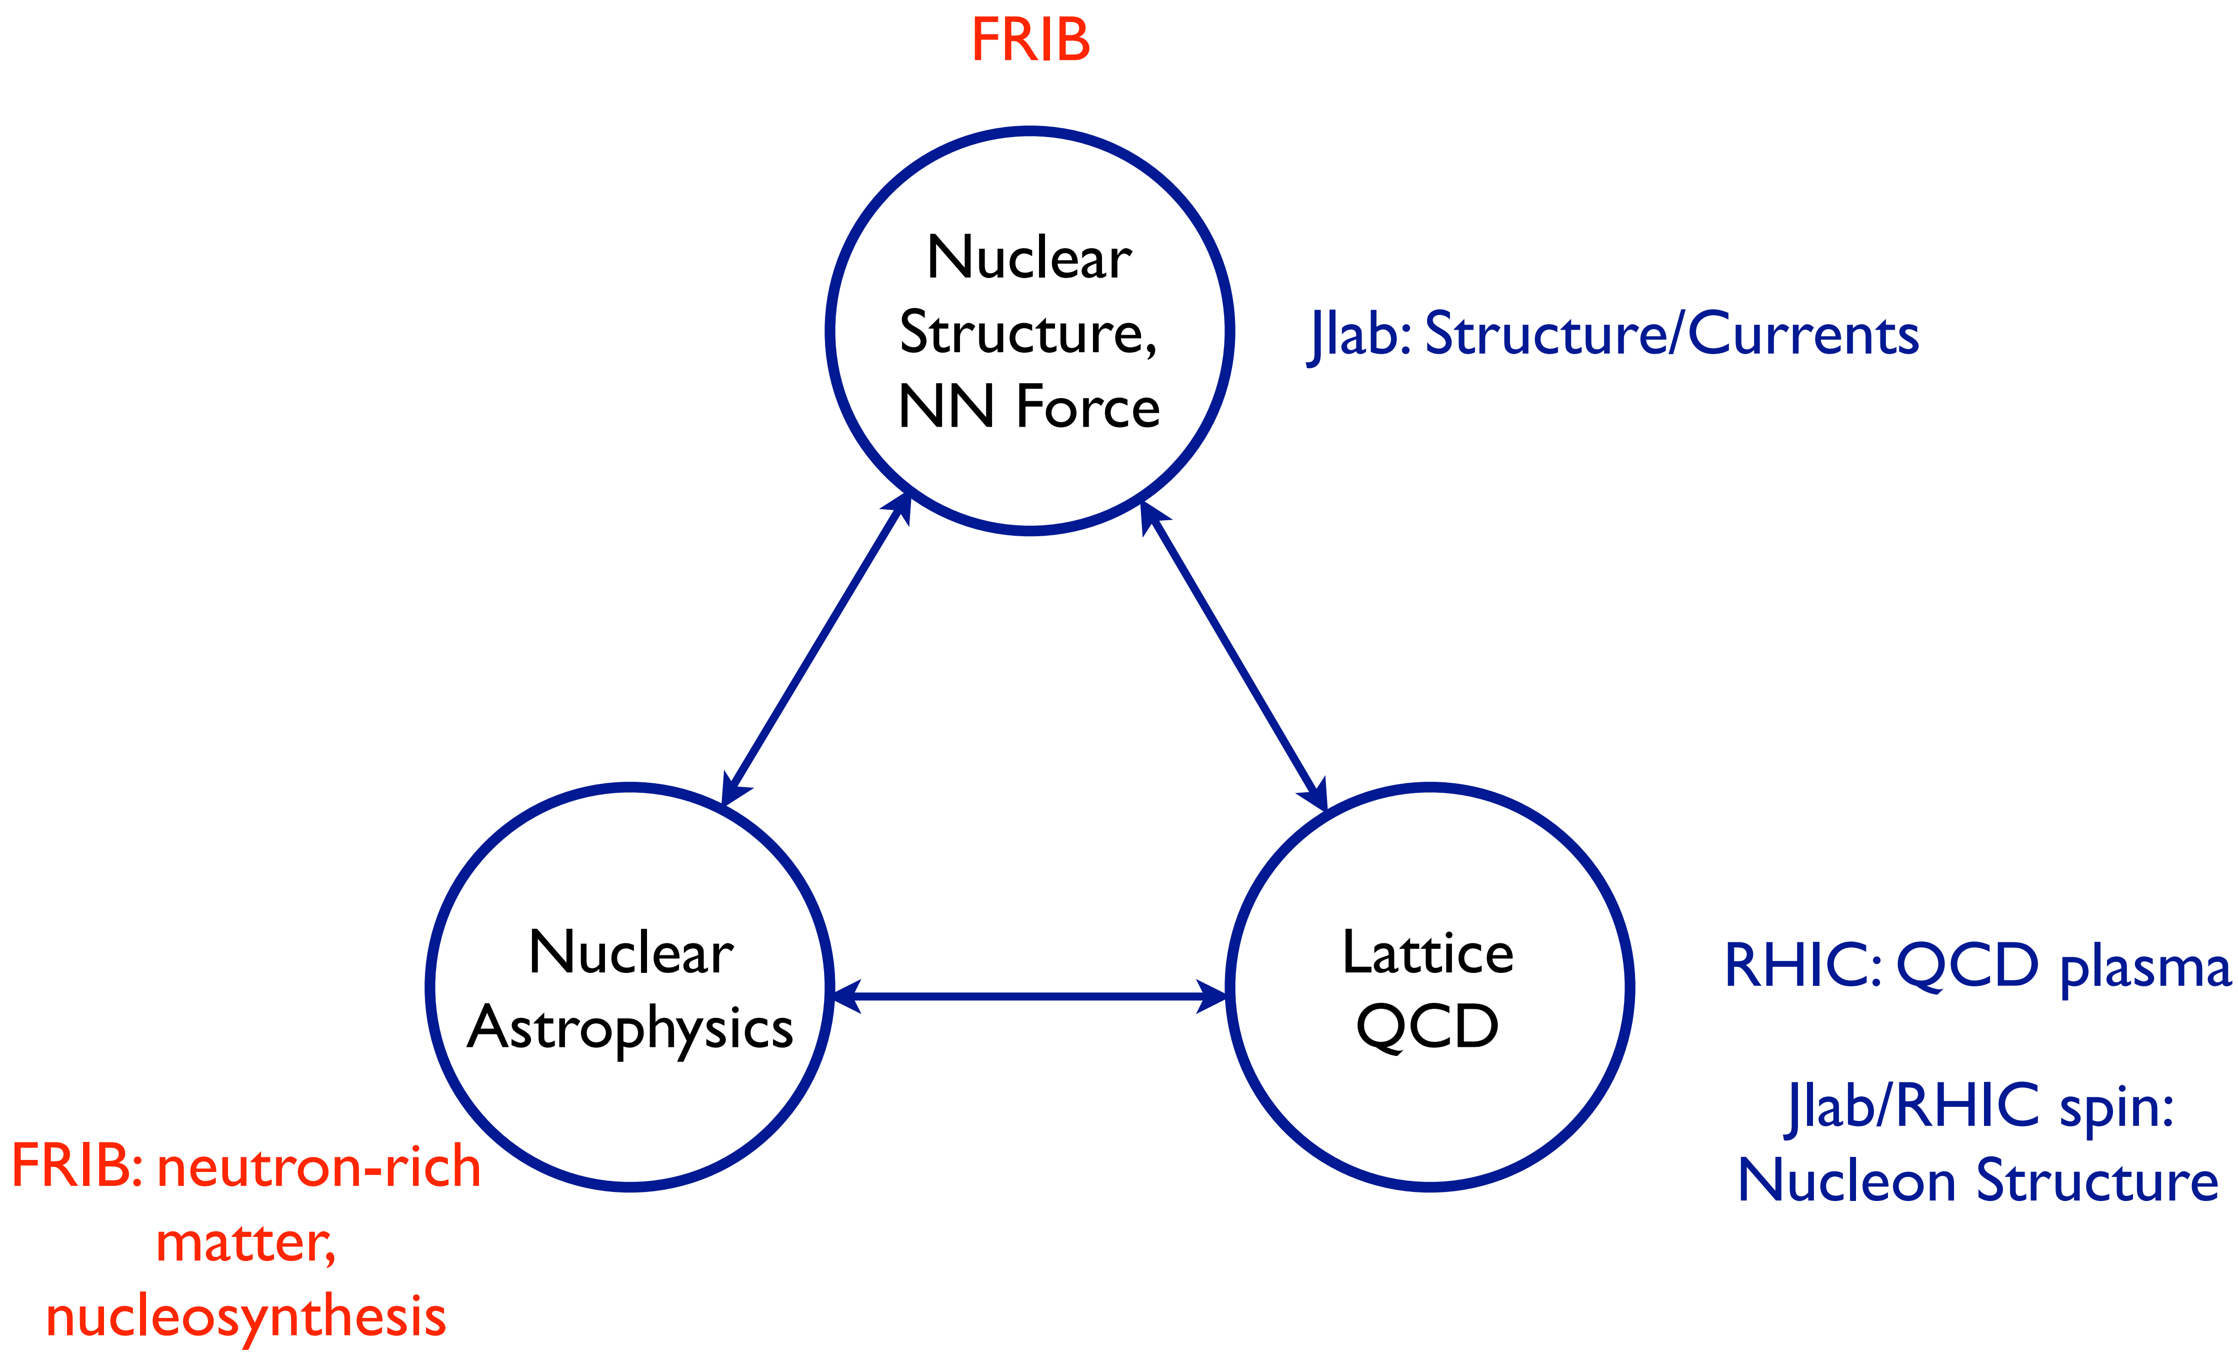
FRIB
Nuclear Structure, NN Force
Jlab: Structure/Currents
Nuclear Astrophysics
Lattice QCD
RHIC: QCD plasma
Jlab/RHIC spin: Nucleon Structure
FRIB: neutron-rich matter, nucleosynthesis

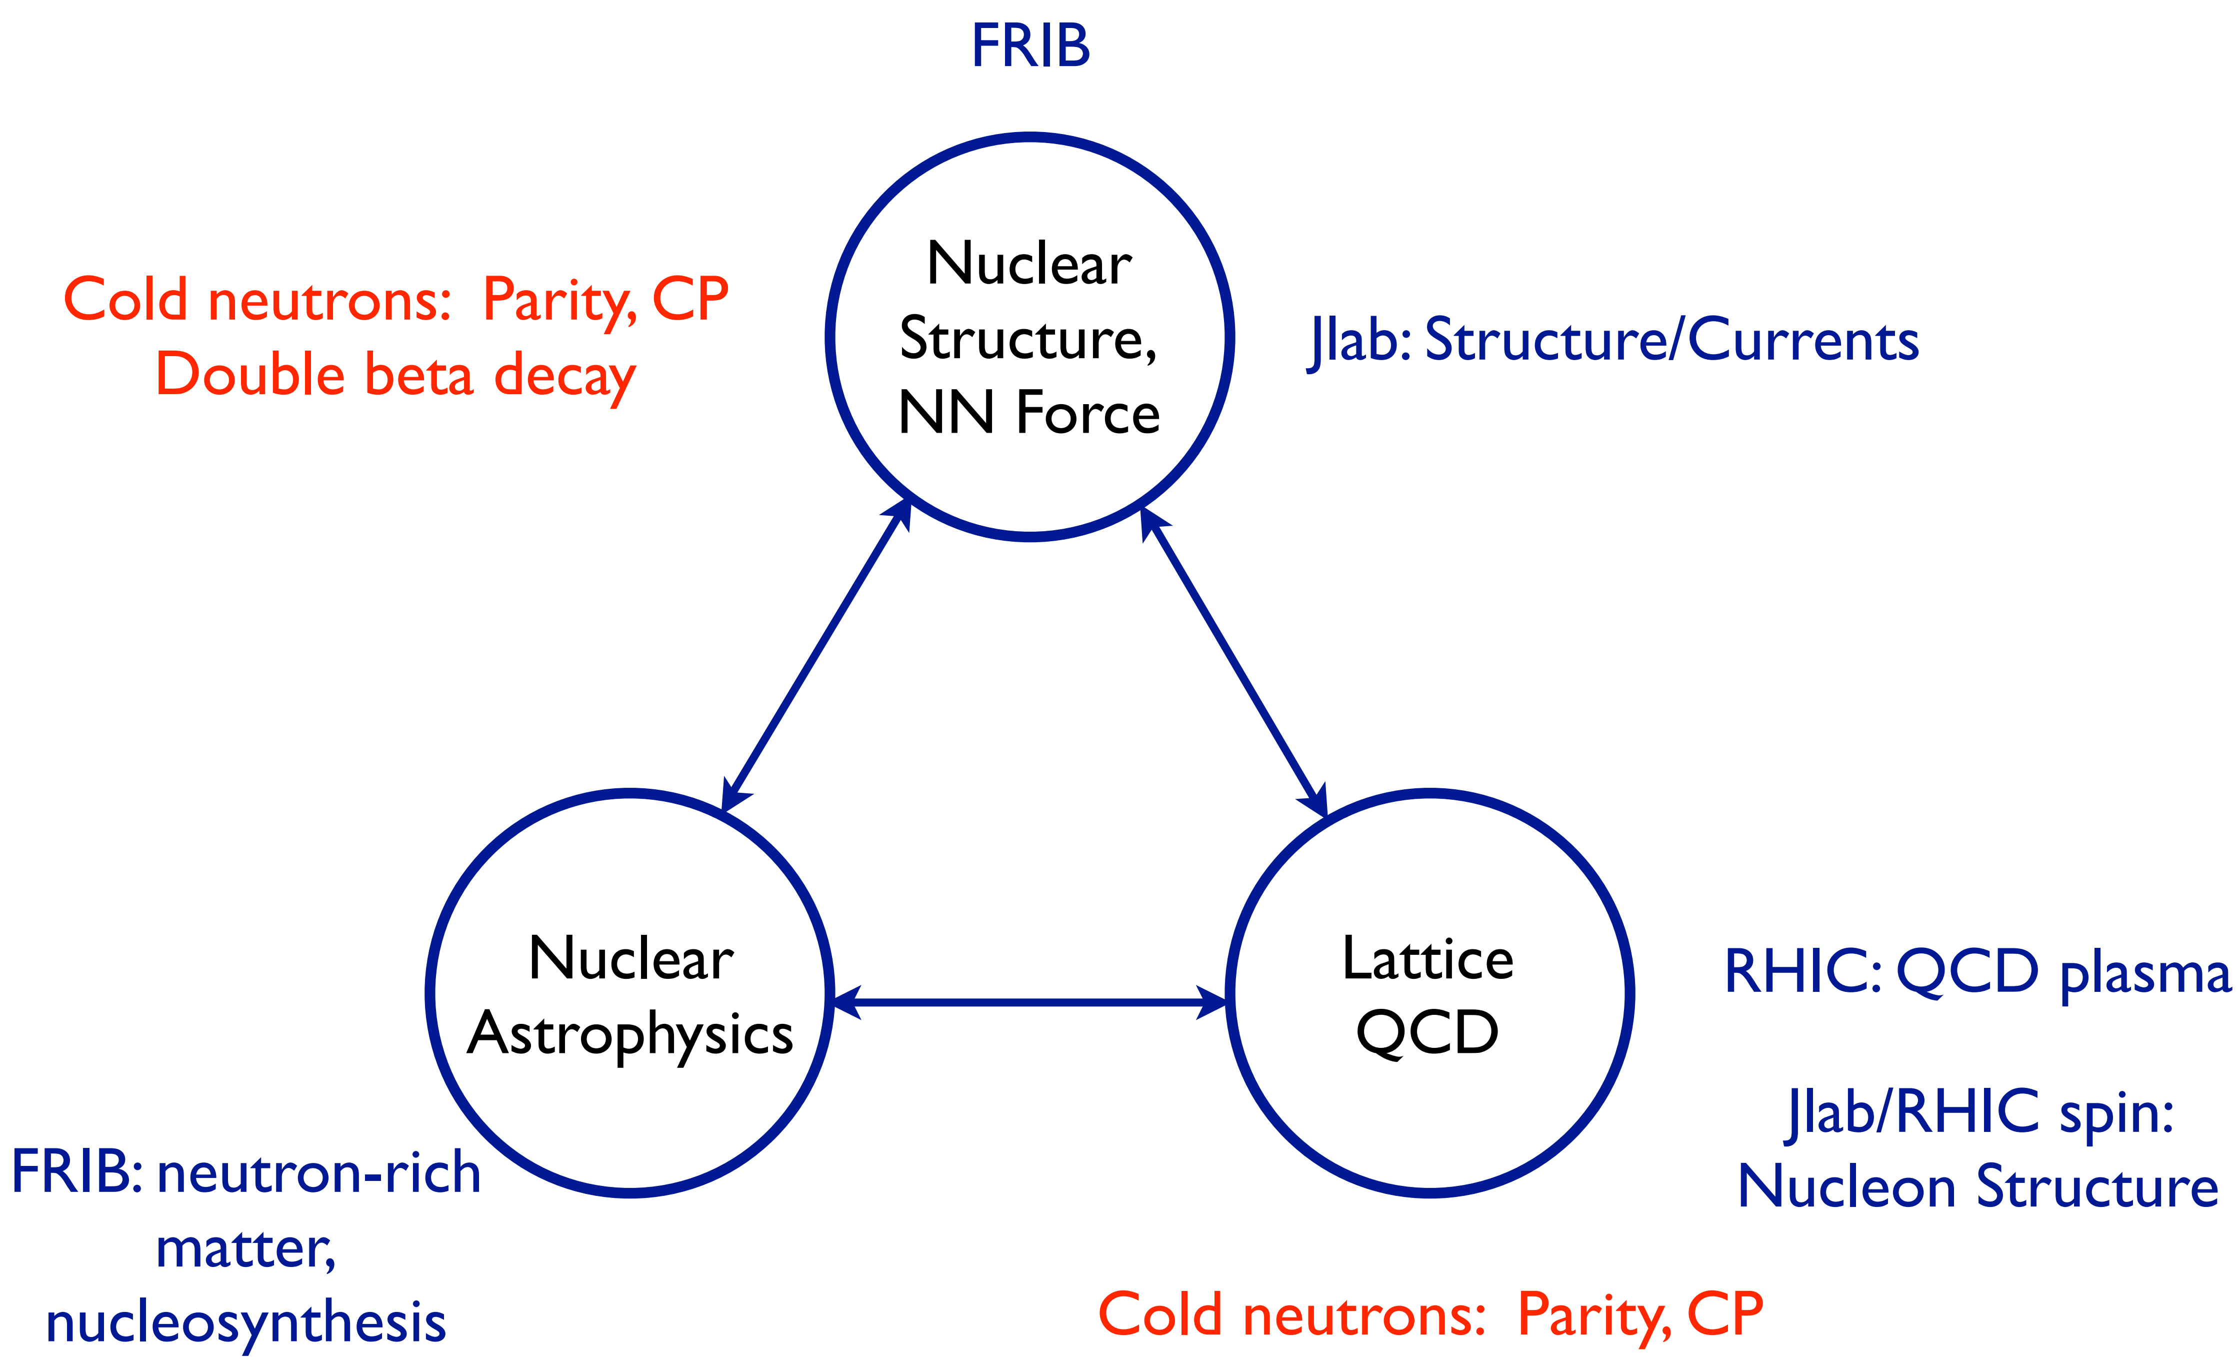
FRIB
Cold neutrons: Parity, CP
Double beta decay
Nuclear Structure, NN Force
Jlab: Structure/Currents
Nuclear Astrophysics
Lattice QCD
RHIC: QCD plasma
Jlab/RHIC spin: Nucleon Structure
FRIB: neutron-rich matter, nucleosynthesis
Cold neutrons: Parity, CP

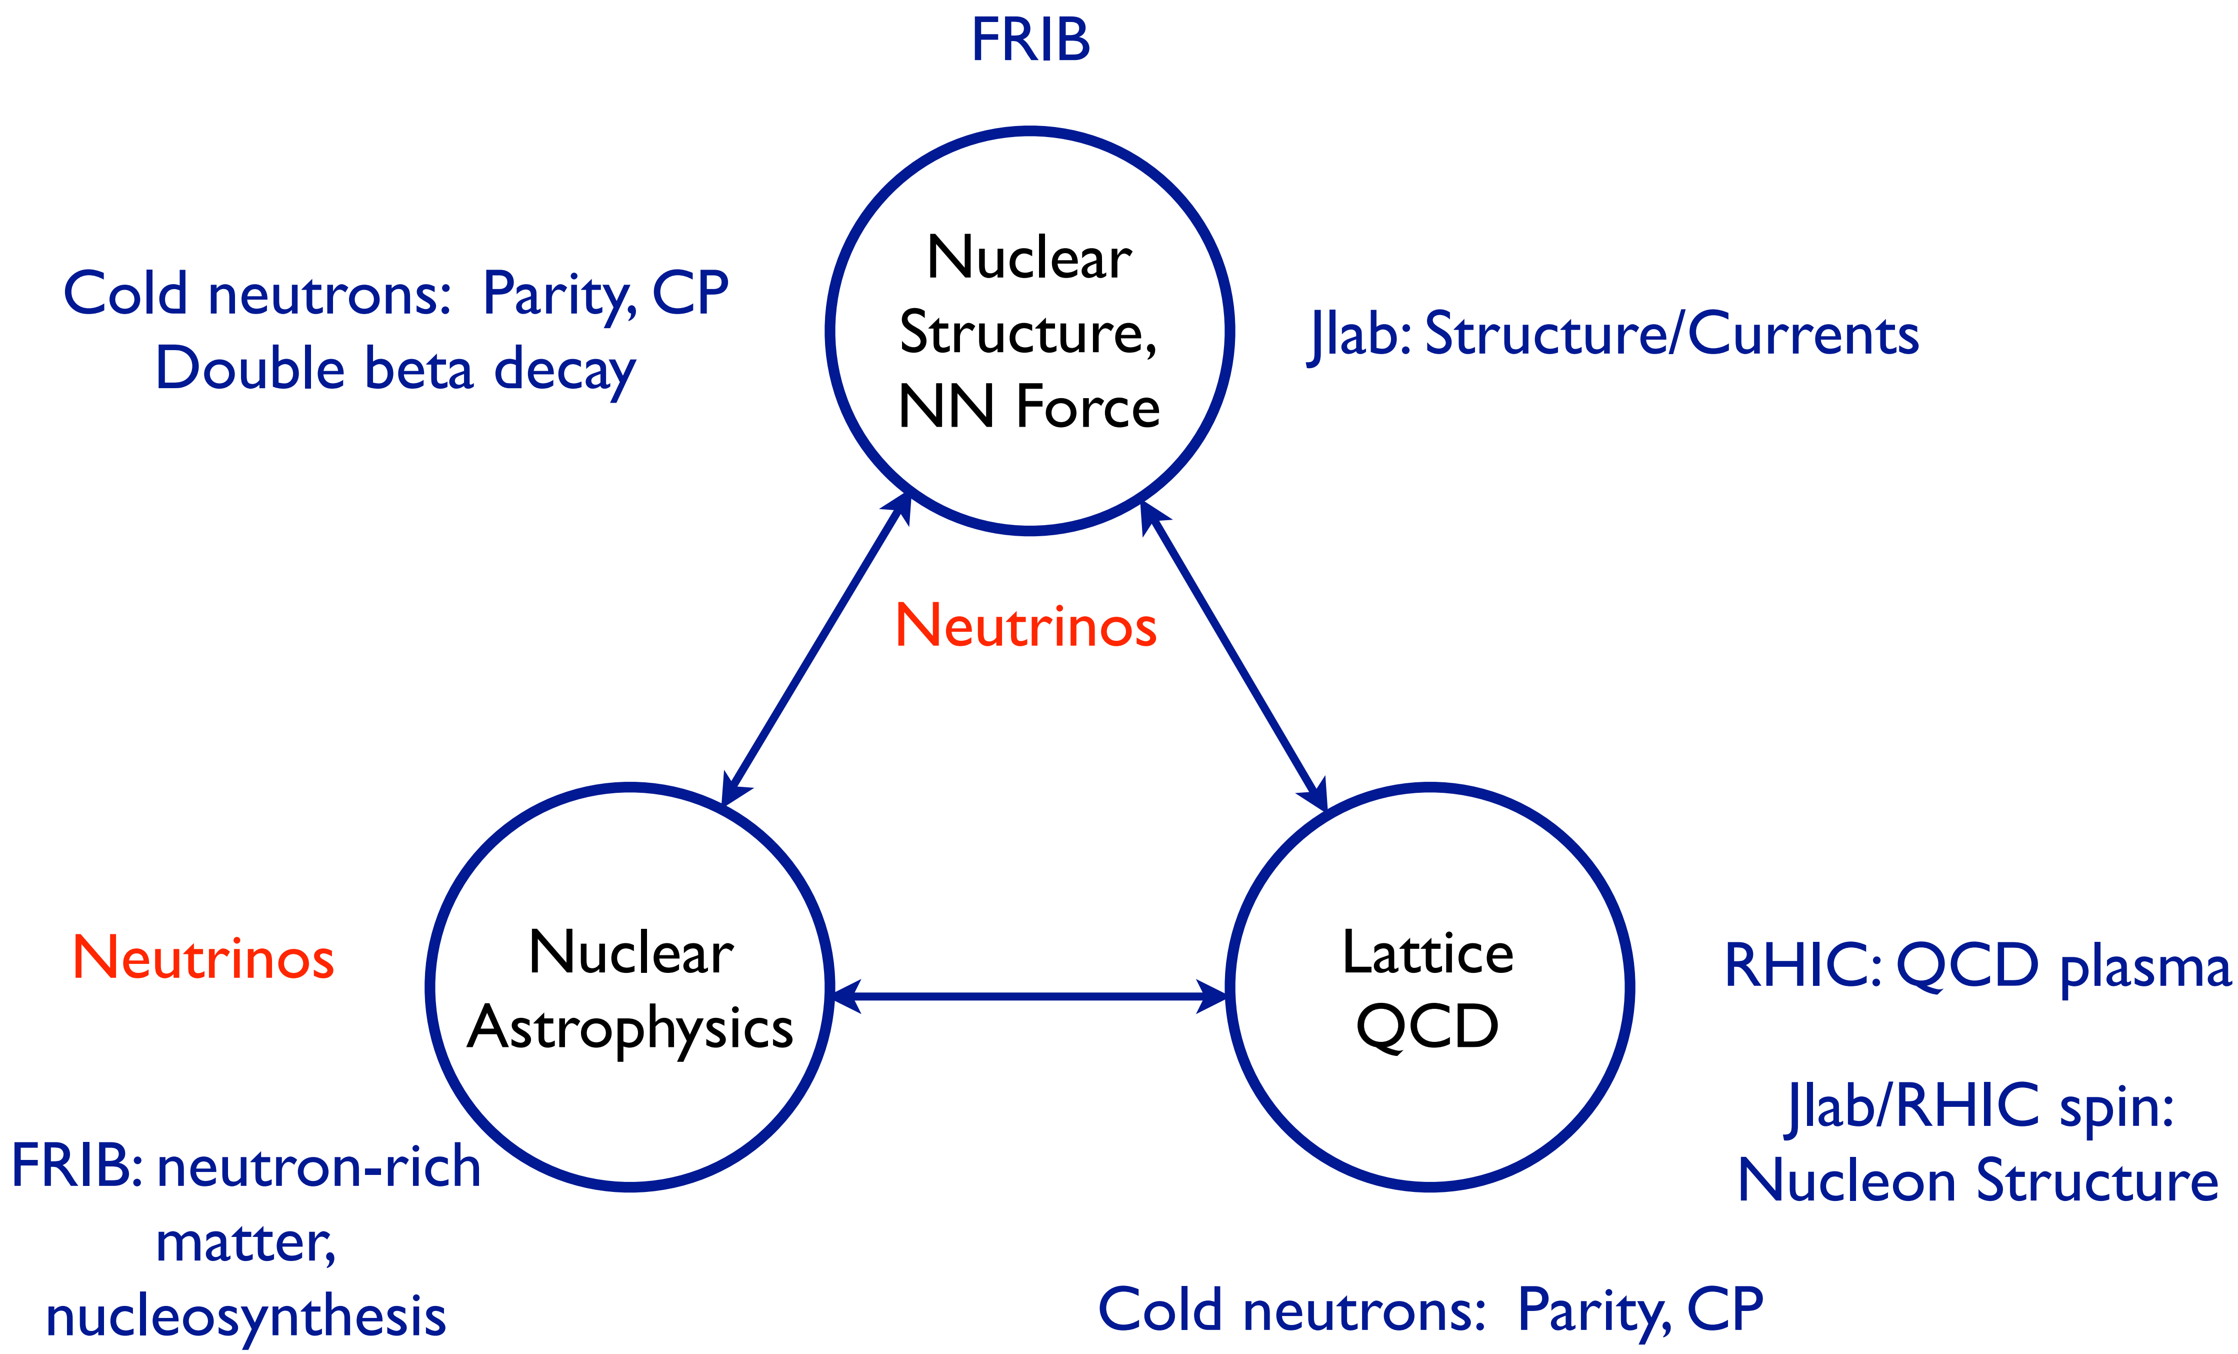
FRIB
Nuclear
Structure,
NN Force
Cold neutrons: Parity, CP
Double beta decay
Jlab: Structure/Currents
Neutrinos
Neutrinos
Nuclear
Astrophysics
Lattice
QCD
RHIC: QCD plasma
Jlab/RHIC spin:
Nucleon Structure
FRIB: neutron-rich
matter,
nucleosynthesis
Cold neutrons: Parity, CP

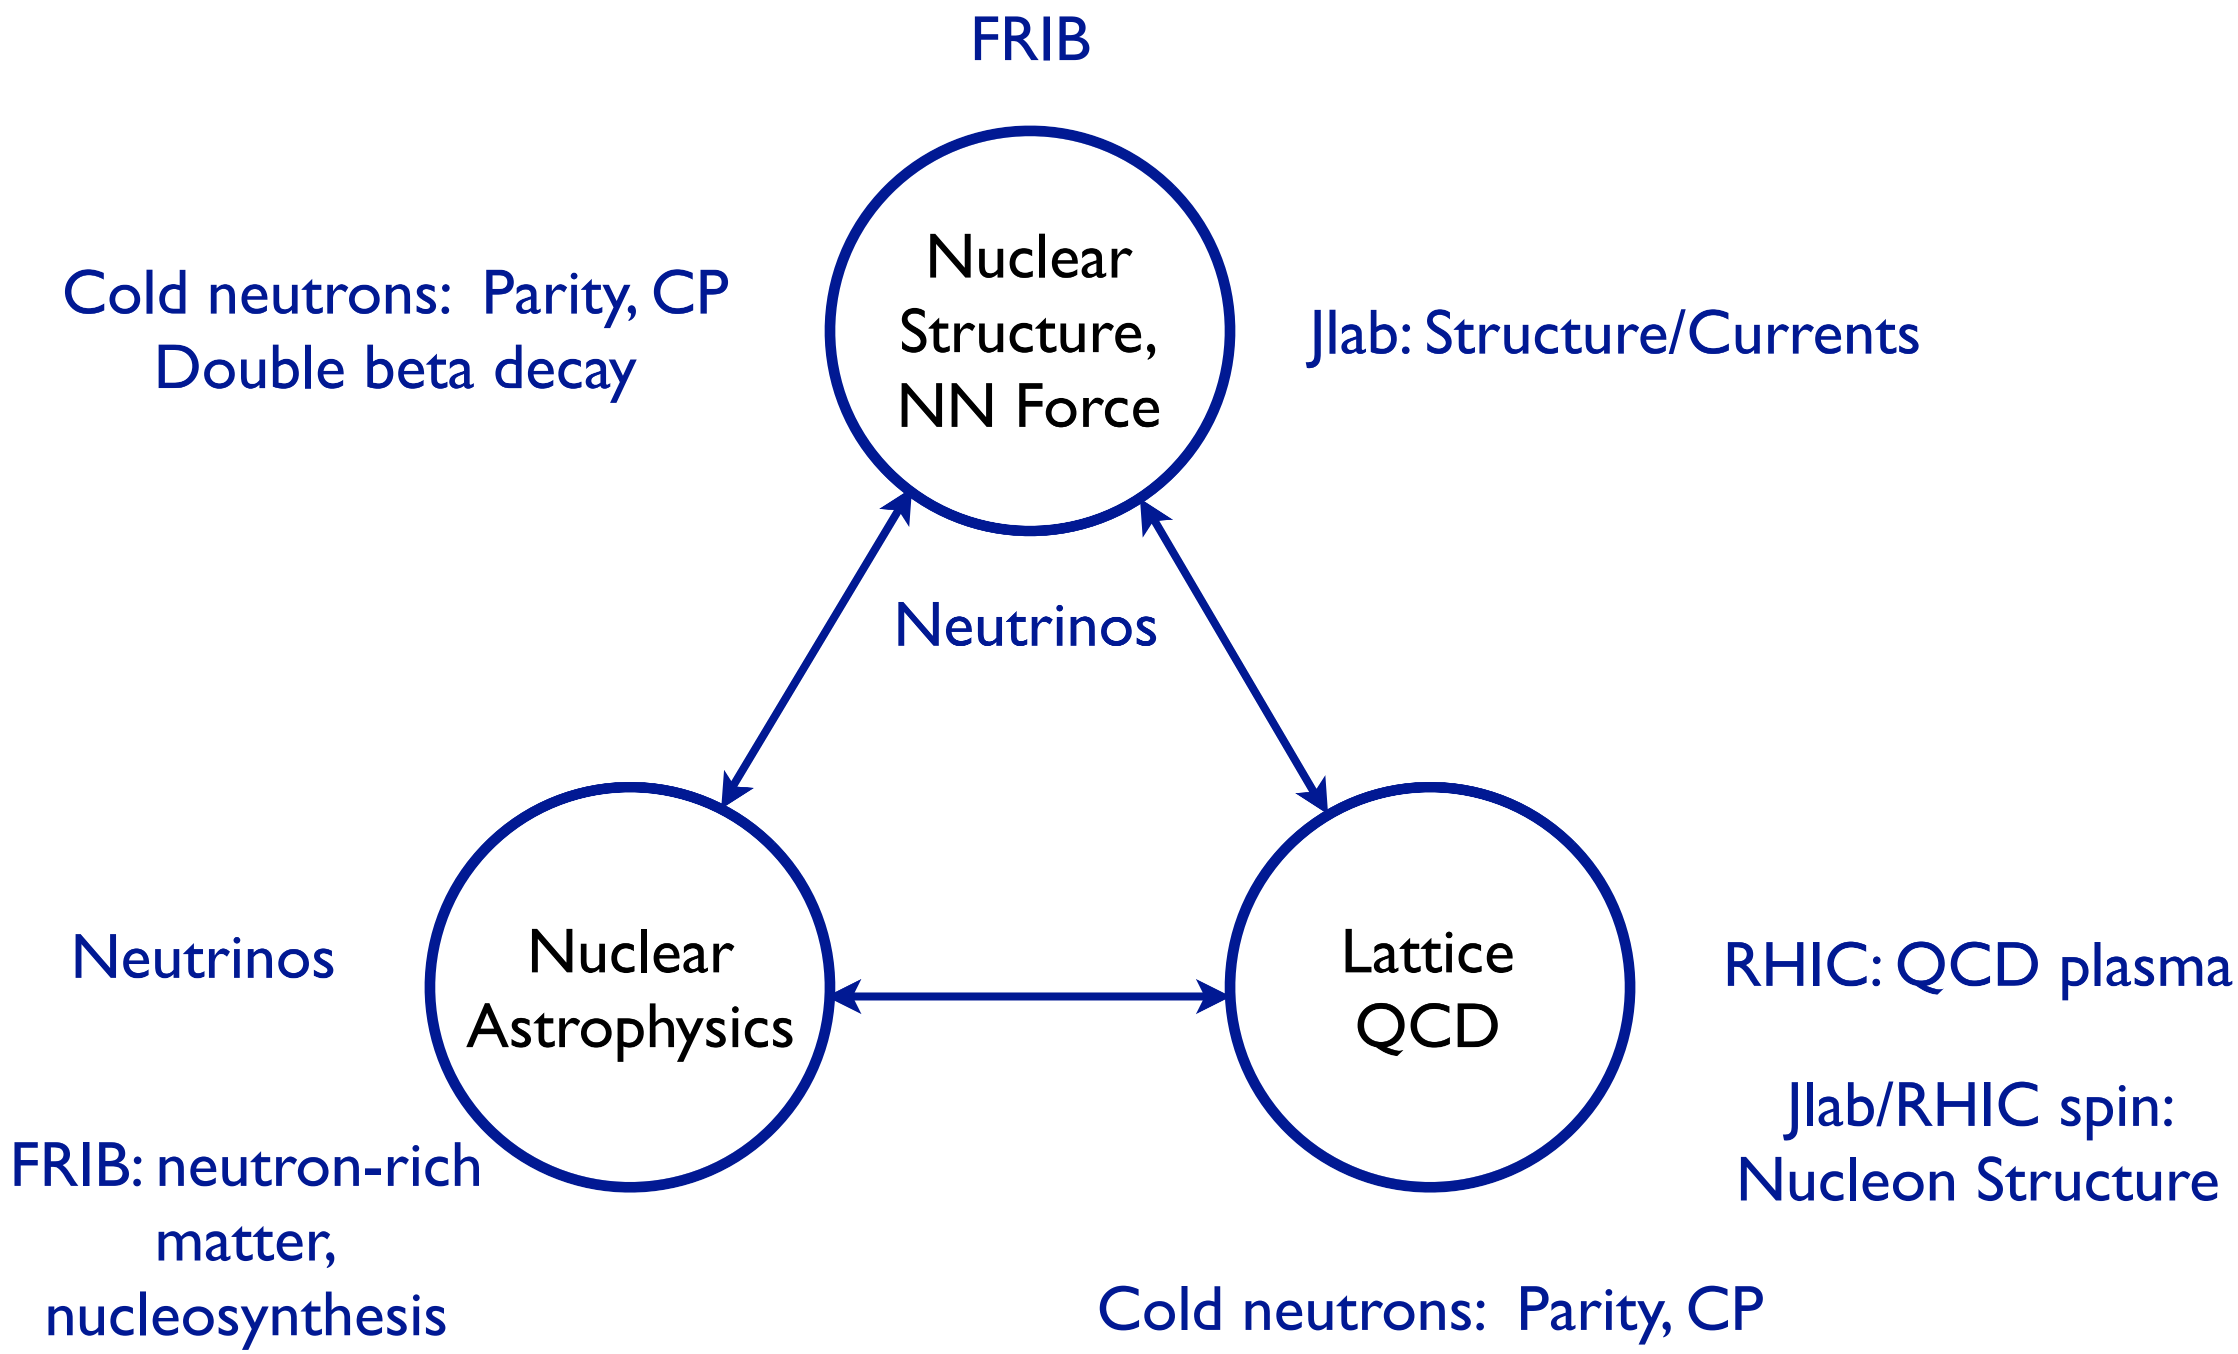

Computational NP not only supports/interprets experiment, but builds the bridges

- lattice QCD ↔ NN potentials ↔ nuclear structure
- neutron-rich nuclei ↔ the r-process path in supernovae

The Case We Must Make within NP

- Computational NP research supports and advances experiment, strengthening the intellectual arguments for new measurements and guiding investment

  Beam time is expensive: it is of great value to have the capacity to predict what can be measured and at what cost in effort

- Our field will be stronger if, at higher levels of the DOE, we are seen taking the lead in advancing Office of Science priorities

  FY14 Senate:
  1) discovery and design of new materials
  2) advancement in biofuels production
  3) the development and deployment of more powerful computing capabilities to take advantage of modeling and simulation … to maintain US competitiveness

- The opportunities arising due to advanced machines requires a new approach to computational NP in which increased investment in teams (e.g., SciDAC) and infrastructure is provided as a separate component of the base NP budget

And Why It Will Be Difficult

- The recommendation we need can be taken from the NRC Report NP 2010

A plan should be developed within the theoretical community and enabled by the appropriate sponsors that permits forefront computing resources to be exploited by nuclear science researchers and establishes the infrastructure and collaborations needed to take advantage of exascale capabilities as they become available.

- And that report also notes why this (and other new research) will be a heavy lift

In recent years the resources that go to the major facilities have made up an increasing fraction of the total budget, especially that devoted to the Department of Energy (DOE) portion. Continuing to increase the share reserved for facilities operations at the expense of the research budget is not sustainable.

The current research budget percentage of 28% is exceptionally low, and with current plans to operate two facilities and construct a third, it will be difficult to reverse recent declines

LRP Strategies?

Approach the Town Meetings emphasizing not what we want, but what we can do to advance TM experimental priorities

<u>Nuclear Structure and Astrophysics 8/21-23</u>

FRIB

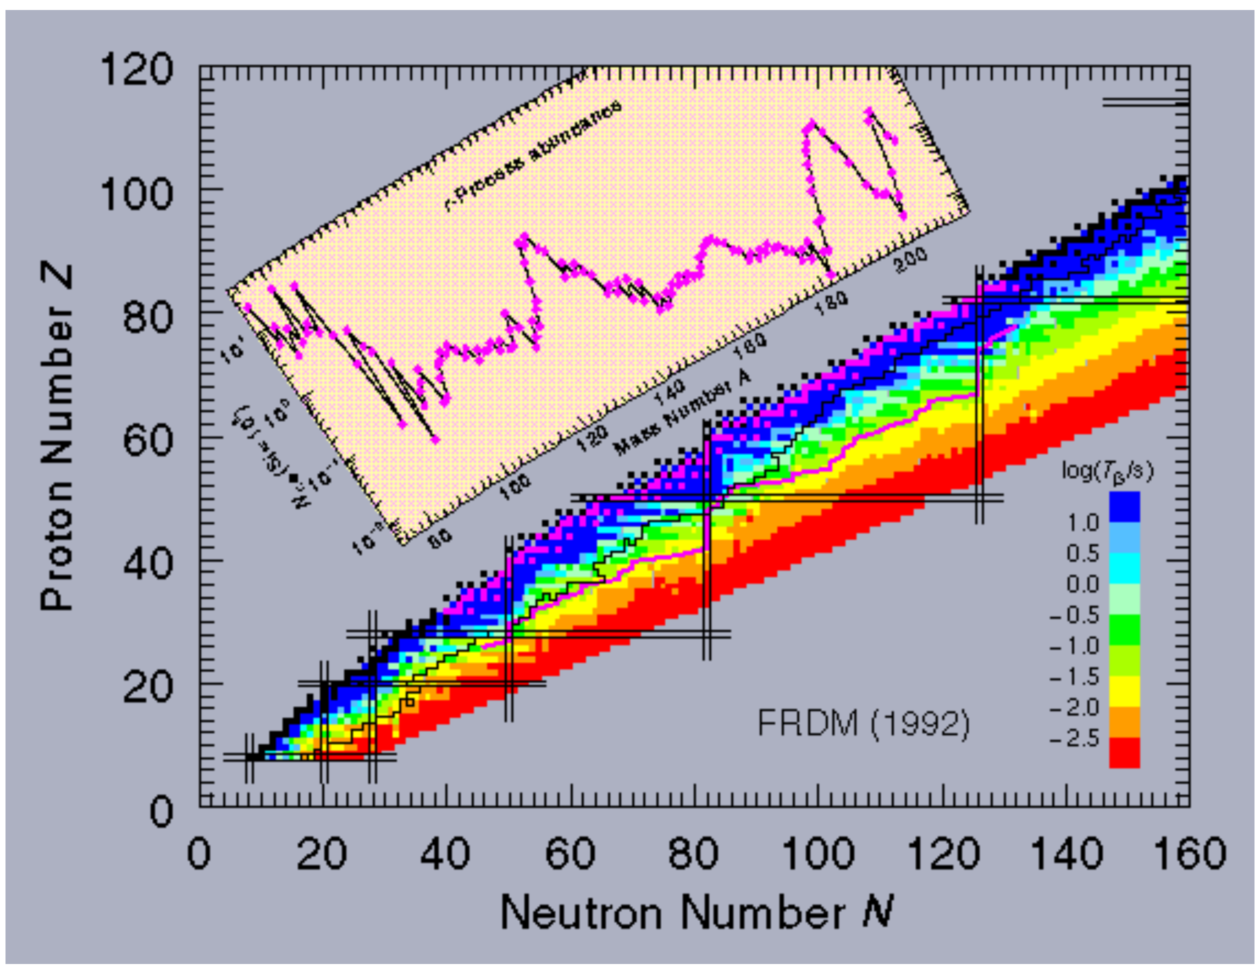

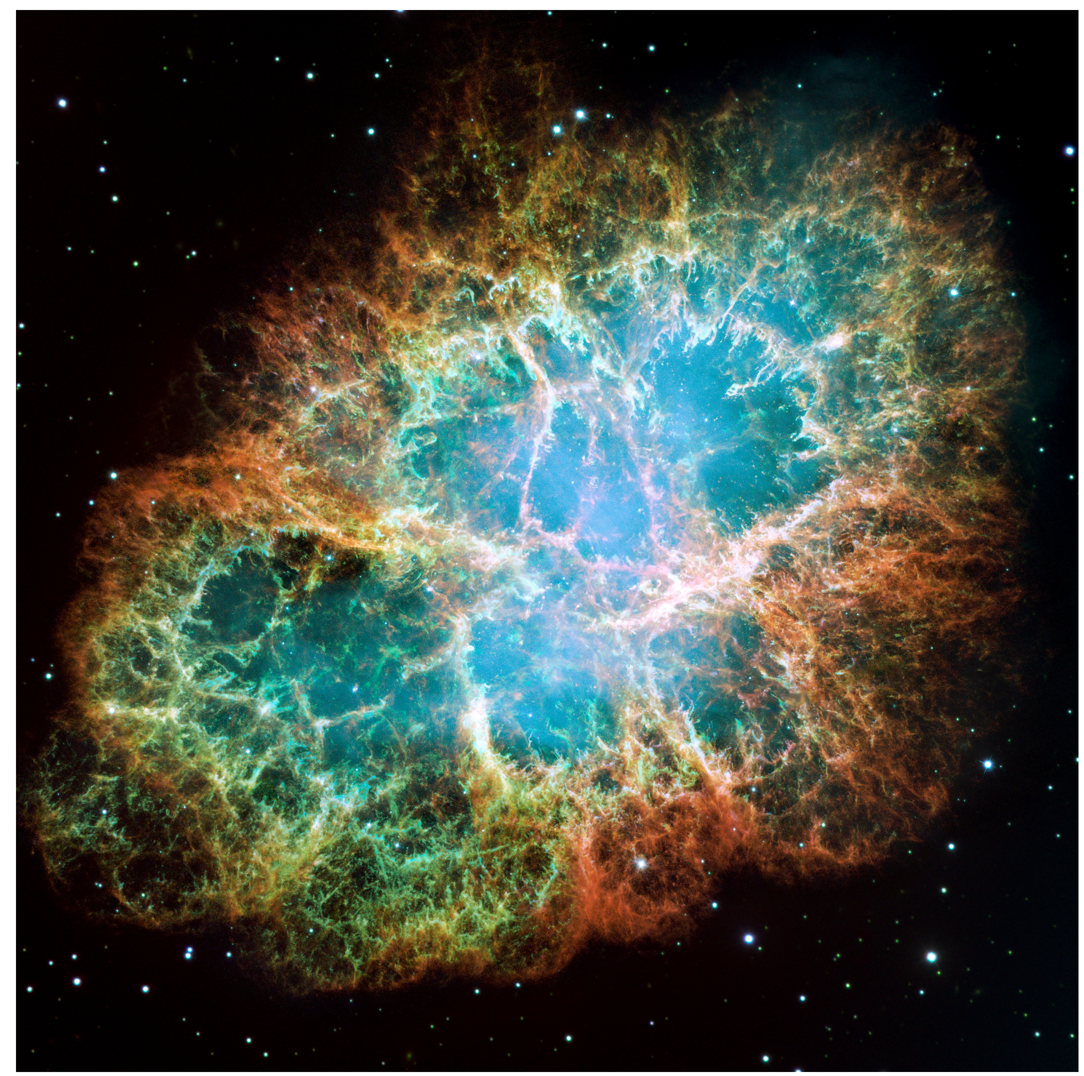

Precise data from FRIB on masses and beta decay lifetimes not impact our understanding of the r-process if large uncertainties remain in defining its astrophysical site(s) — the path and thermodynamic conditions under which heavy nuclei are synthesized

Yet investment in, e.g., SNII simulations has decreased, despite the daunting challenge of realistic 3D simulations

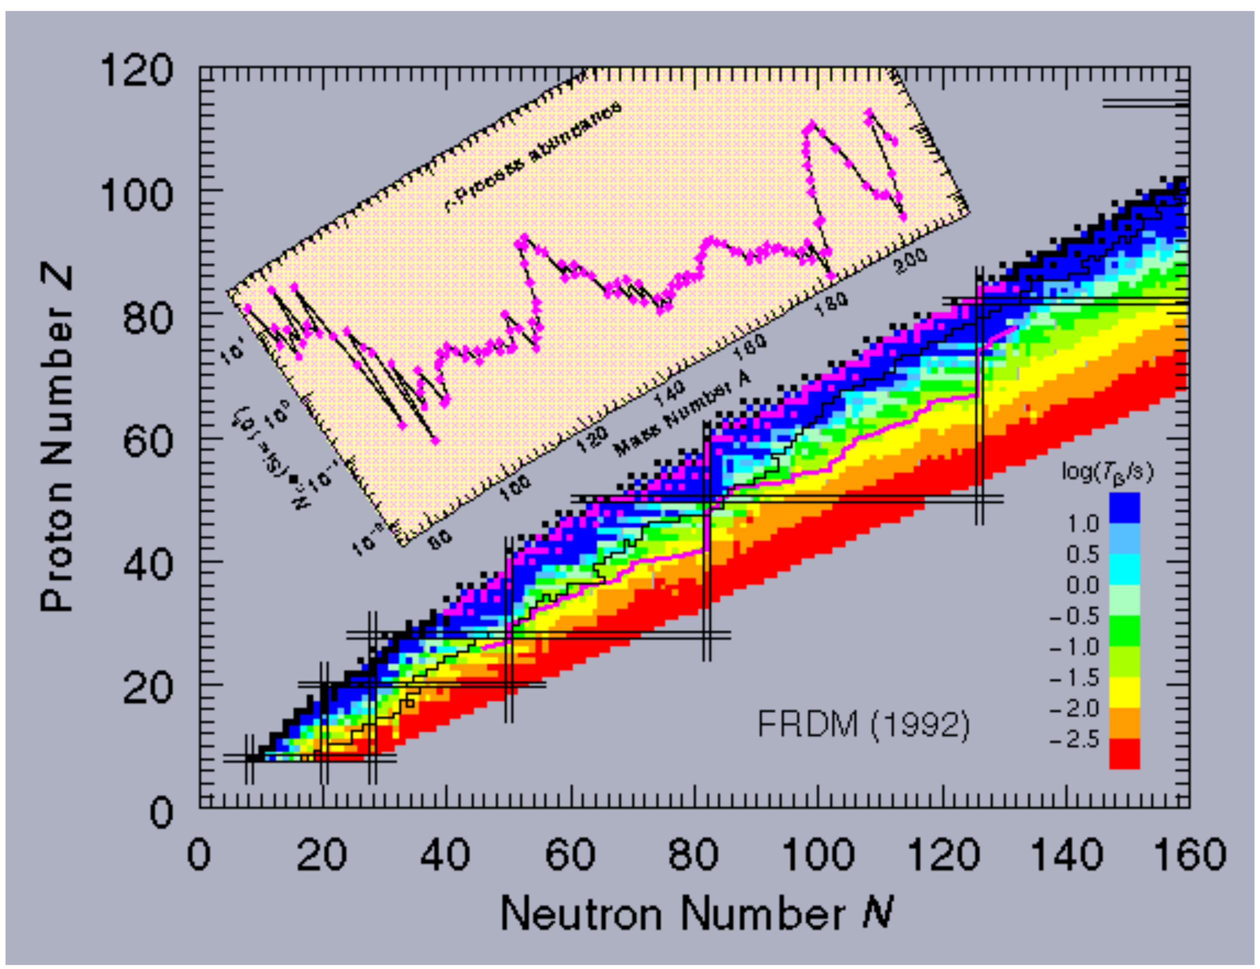

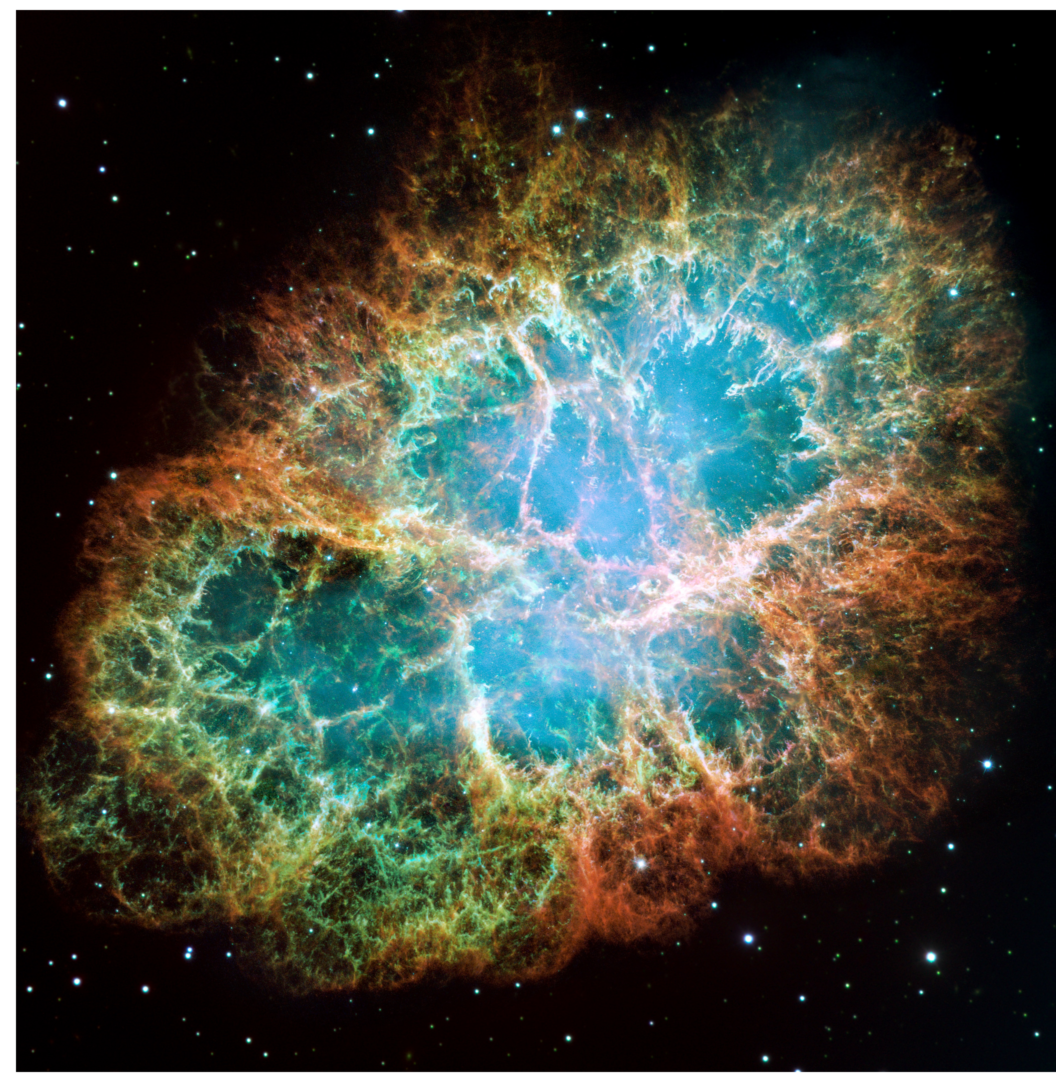

<u>Fundamental Symmetries, Neutrinos, Neutrons, and Nuclear Astrophysics 9/28-29</u>

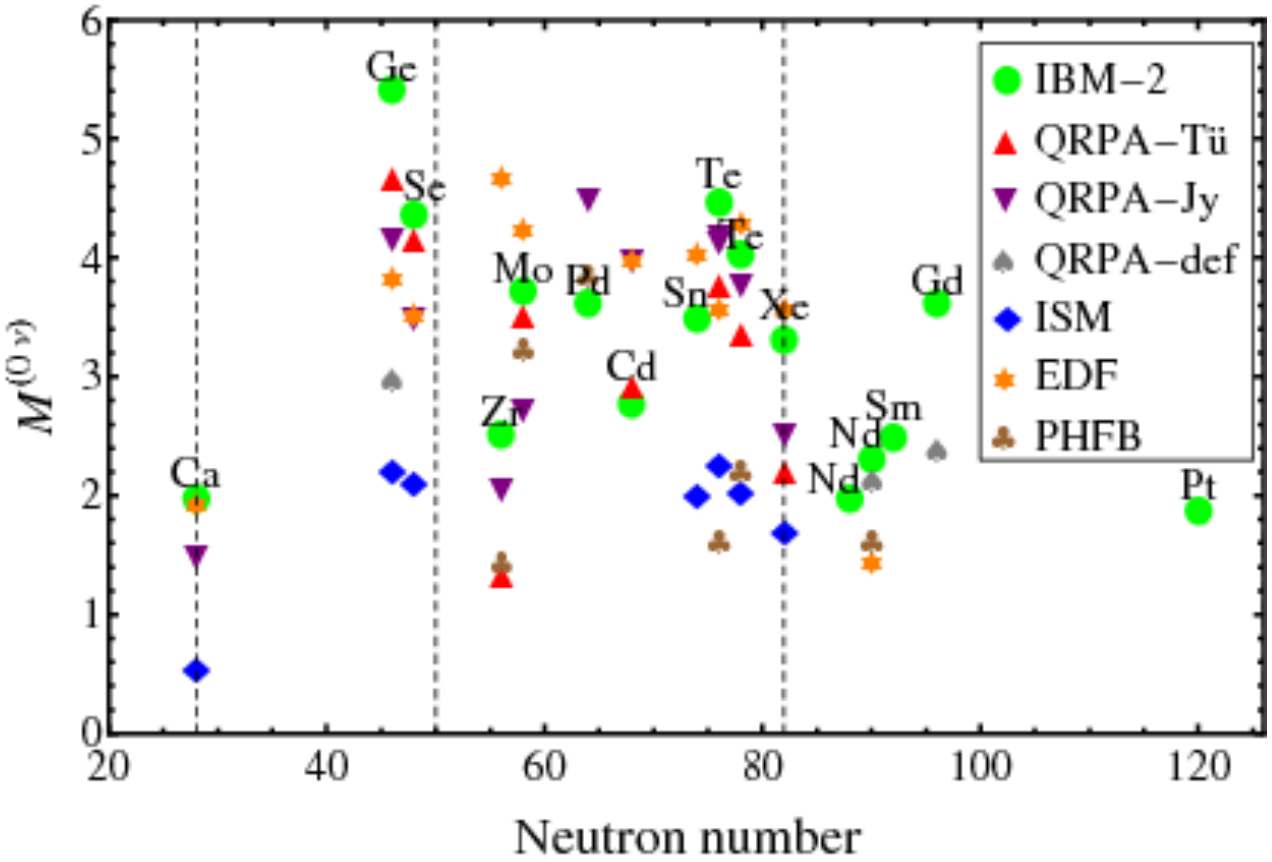

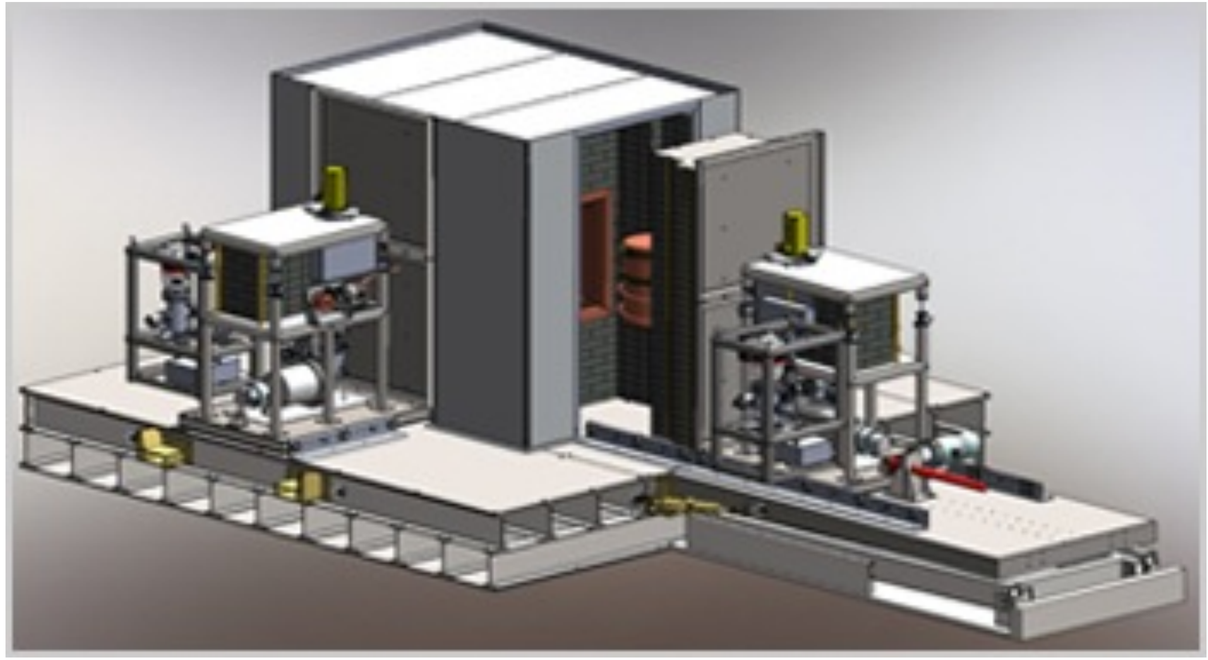

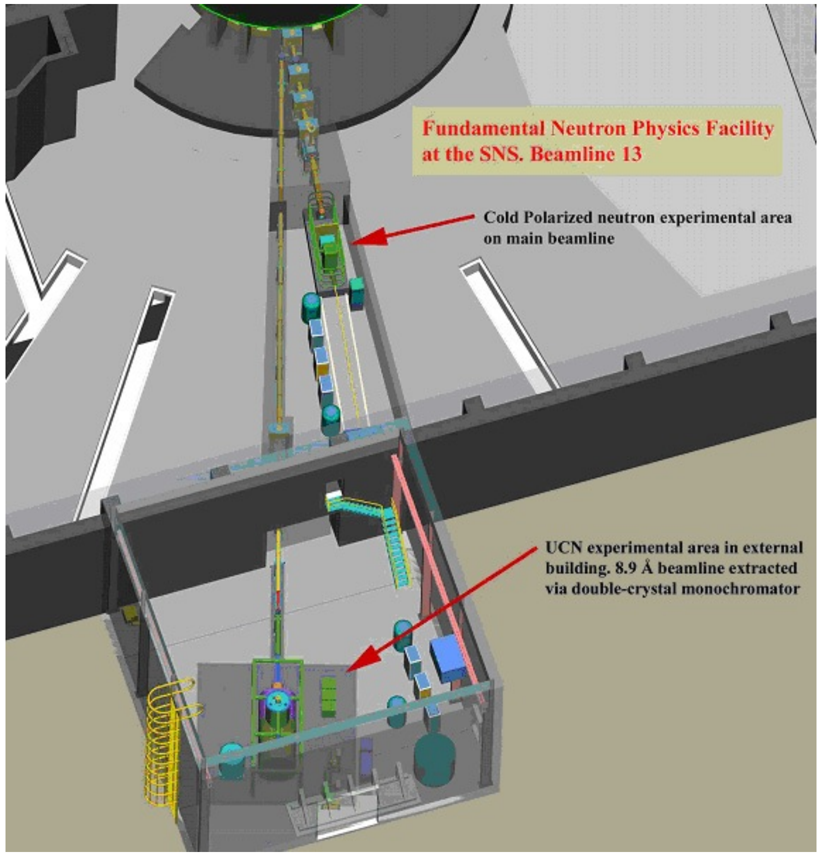

This TM's priorities are likely to be double beta decay, nedm, and neutrinos. Desperate need for theory help from nuclear structure, LQCD:

- effective operators for double beta decay
- LQCD relationships between nedm limits and underlying parameters, like θ_{QCD}
- neutrino response functions for a variety of next-generation oscillation experiments

<u>Hadron and Heavy Ion QCD 9/13-15</u>

A large set of LQCD — hot and cold — and nuclear structure challenges that must be addressed to support the experimental programs, but including

- hybrids and exotics to be explored at JLab-12 GeV
- the hot QCD EoS
- the EIC program — nucleon spin, quark and gluon energy loss, gluon saturation

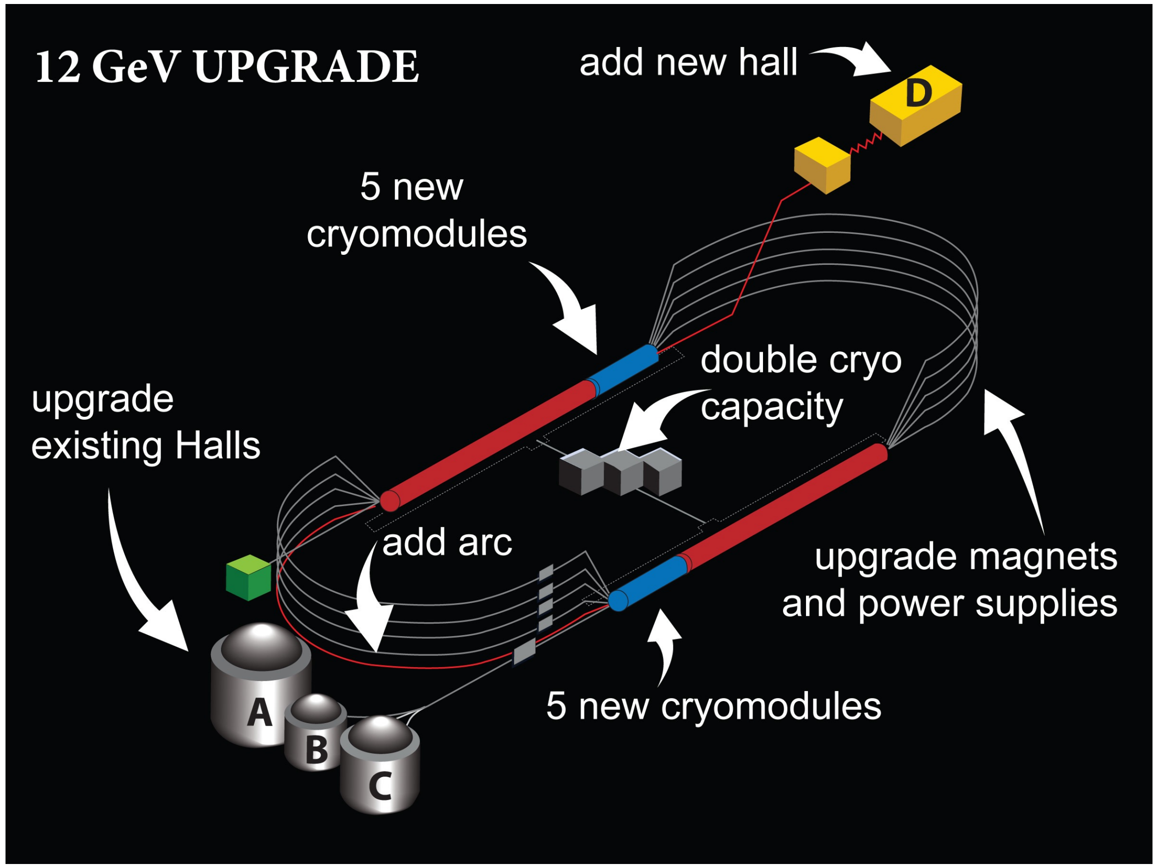

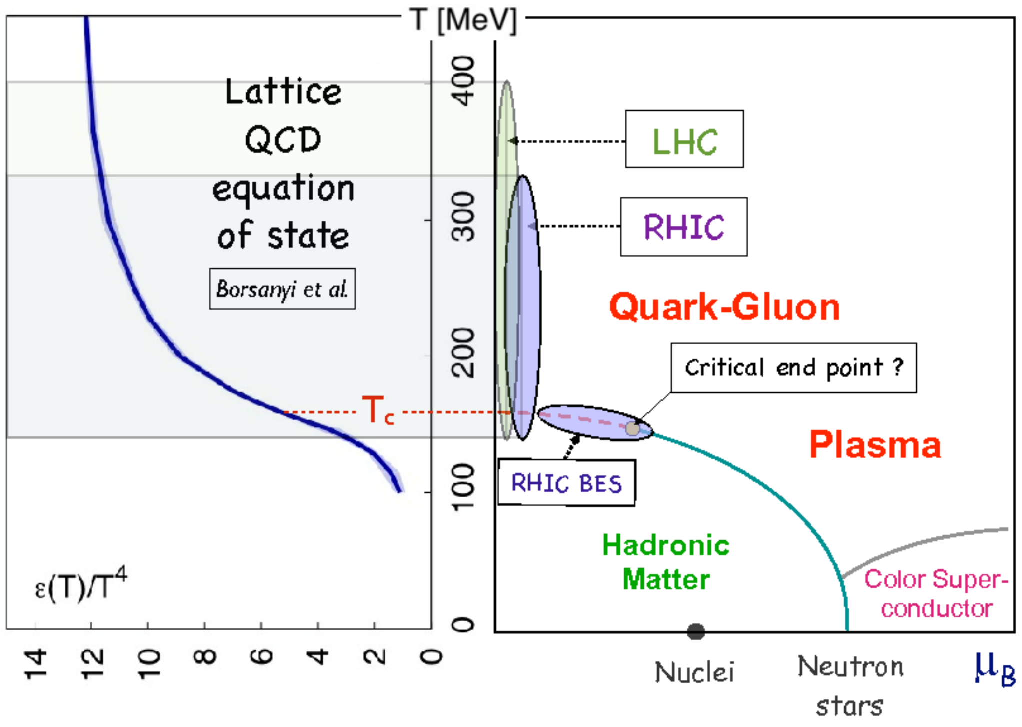

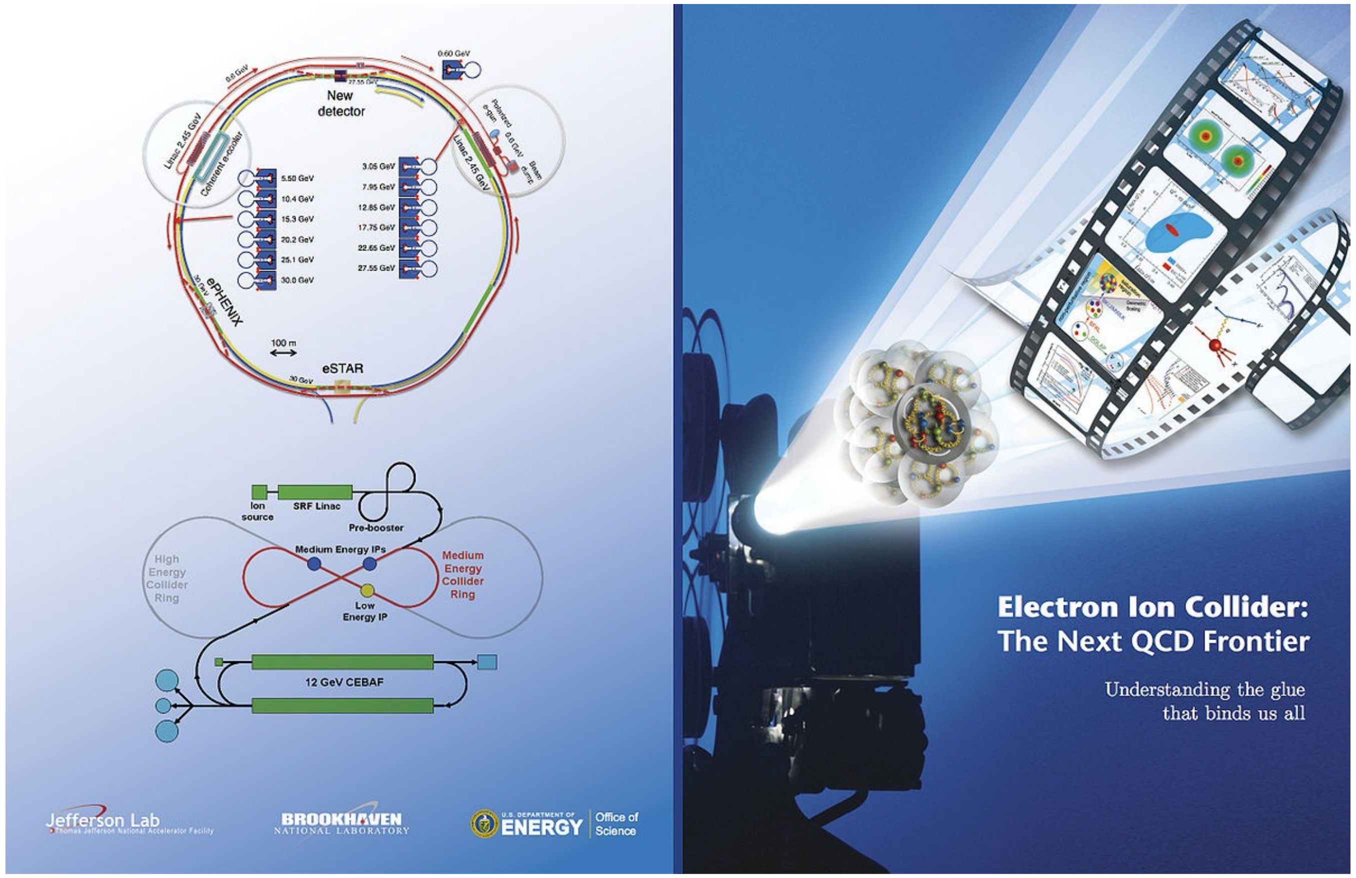

Education and Innovation 8/6-8

Given the DOE expression of support for computing because of its relevance to national needs — economic, security — our ability to train scientists who are also adept at simulation is an important plus. Interesting programs are being created

e.g., Berkeley Designated Emphasis in Computational Science and Engineering

"The dramatic increase in computational power for mathematical modeling and simulation has led to the fact that scientific computing now plays a significant role in the analysis of complex physical systems, such as computer chip manufacturing, battery modeling, turbine design, aircraft prototype testing, climate change and star formation, to name a few. More recently, too much data has become another compelling problem: radio telescopes, DNA sequencers, particle accelerators, sensor networks, social networks and the internet all collect much more data than humans can analyze and understand. In this case one needs to use statistics and machine learning, along with data visualization, to extract useful information from the data. The solutions to the mentioned problems share many mathematical, statistical and computational techniques in common.

The mission of the CSE program is to help educate UC Berkeley students in these common techniques, to help them solve CSE problems across a wide range of disciplines. The CSE program is committed to the development of new curricula and expanded programs aimed at development and propagation of the use of numerical and computational tools to further research across multiple disciplines. To that end, the CSE program will actively support the training and multidisciplinary education of scientists, engineers and technical specialists who are experts in relevant areas."

I. We need to decide on a TM/Resolution Meeting strategy: A straw-man proposal

Town Meetings:

- Various of us being active in the TMs, connecting the high-priority experimental initiatives with specific calculations we will be able to do in this decade

- Consequently, argue for the inclusion of simulation and computation in the key experimental recommendations
  - — there is some empirical evidence that this is a better strategy than fighting for individual theory/computation bullets

- A "5% for computation" strategy: theory is under great stress and does not have the capacity to support the kinds of teams we need to make progress on the fore-front problems. View simulation as a component of new facilities, a tool that can guide experimental programs, setting benchmarks
  - — that is, position simulation/computation as a component of new experimental initiatives, not as a component of theory
  - — support others who will argue for increasing the research budget

We have a precedent: the theory add-on that led to the INT

At the Resolution Meeting

- Come prepared with a report from this meeting that makes a case for a coherent initiative of plausible size
  - — the detailed physics is less important than the inspirational big picture: "why now, why us?"
  - — an executive summary that gives the big picture, with three or so "killer apts" taken from our TM successes (r-process, double beta decay, JLab-12 GeV, EIC)

- Have our "talking points" simple and clear, so our message is heard:
  - — the coherent support of the Town Meetings for an initiative in computation and simulation aligned with and a component of NP's new projects
  - — the importance of NP leadership, so that our field is seen as aligned with a key OoS priority
  - — now is a special time: 300 petaflops of computing to become available at three labs where NP is embedded — now is the time for the field to help us increase our share of this pie

II. We need to equip DOE NP (and NSF Nuclear) to argue for new resources for us

- The science case — made well already, but can be reinforced in the TM process

- Paying attention to DOE — how it approaches computing, and what it values
    - — partnerships with ASCR: NP is woefully under-invested in SciDAC
    - — use and exploitation of the most advanced machines
        - a focus ought to be CORAL: the ORNL/ANL/LLNL collaboration to deliver three 100pflop platforms in 2017/18
        - groups like USQCD have the sophistication to quickly move to new machines

    So I think our plan has to communicate that we want to be leaders in the use of next-generation machines, and we need to start building the collaborations and infrastructure now if we are to be ready in 2018

- Innovative workforce ideas — equipping grad students, postdocs to deal with computational science circa 2020

“... the fiercely competitive global landscape with China’s goal to educate one million students in high performance computing.”

“Currently, the U.S. has three of the five most powerful computers in the world, but our global competitors are not far behind. Maintaining a cutting-edge, domestic advanced computing capability, however, is a crucial component to achieving our mission and furthering the science that underpins advances in energy technology, environmental remediation, and nuclear stewardship. This capability requires both advanced hardware and, equally important, the advanced software, algorithms, and operating systems that are optimized to take full advantage of our investments in new machines”

FY14 Budget Testimony: E.Moniz

The boss has a very realistic view of what is needed — so if we help by providing DOE NP with a good plan, we might receive a sympathetic hearing

III. Innovative workforce development

- My own experience recently has convinced me of the value of students and postdocs that are equally adept at physics and computer science

- I believe if we made a well-considered proposal to NERSC/ORNL/ANL to form a partnership to train students and postdocs jointly in physics and computation, we would get a positive response
  - — the students and postdocs could reside in our groups, but spend significant time at the lead facilities
  - — this would start building university-lab connections that could have many other positives
  - — the local Berkeley version of this involves a joint mentorship of PhD students, and a designation on their PhDs — a nice model

- ASCR and the labs have, on their own, started a new program of postdoctoral fellowships "in the area of development of computational algorithms and their application to scientific areas that are of interest to DOE"
  - — can we get on board?

Formulate a specific request, the recommendations in our white paper:
a balance between our priorities and what might be possible

#1 Address the under-funding of SciDAC: the NP shortfall was estimated this past August as approximately $2.5M/year

- — this is something DOE NP can do on its own!
- — we may want to think about how to quote the number (e.g., percentage increase)

#2 Address our people needs: if we ask for more postdocs for all of thus, this will be a nonstarter, as research across the field is under stress

- — a nationally focused "young researcher ⟷ interdisciplinary training" program that involves ASCR and builds a partnership between our distributed NP groups and NERSC and the CORAL partners strikes me as the most likely avenue to success (Yelick)
- — encourage the TM on Education to echo the Tribble Report:

"People remain the key factor. In particular, early-career scientists working at the interface between nuclear theory, computer science, and applied mathematics are critical to make future impact, especially in the era of extreme computing that demands the novel coding paradigms and algorithmic developments required by novel architectures."

<u>#3 Address our need for cycles:</u> the cycles will exist of the CORAL machines and on NERSC's Cori, all available in the 2016-18 time frame

- — our job is get a greater share of the pie
- — these are the machines that the DOE bosses identify with national needs
- — we increase our competitiveness for cycles by
  - associating NP with innovative programs (like #2) that advance ASCR and lab goals in addition to our own
  - getting our field to stand out in terms of its investment in collaborative computational science (#1)
  - grass-roots efforts to build collaborations with ASCR CS/AM partners and with lab colleagues

Bottom line: keep in mind Mick Jagger and the INT

- enormous effort led to a $1.7M project in 1990 with a five-year ramp
- scaling that to today is plausibly a $5M project with a five-year ramp — nicely congruent with the big machines coming, SCIDAC IV
- that is enough to make NP SciDAC healthy and create a targeted program for young researchers, the latter hopeful resigned in partnership with ASCR and including the goal of better integrating NP into CORAL and NERSC
- camel's nose theory: make this envelope and an NSAC subcommittee the LRP "product," then work with George and Ted