Leadership-Class Configuration Interaction (LCCI) Code/Environment Developments

Report to Annual SciDAC/UNEDF Meeting
Pack Forest, June 22-26, 2009

Based, in large measure, on
Report of LCCI Workshop goals/accomplishments
San Diego State University
March 12-14, 2009
* Posted to UNEDF website*

James P. Vary
Iowa State University

Who are we?

Self-identified & Self-organized UNEDF subgroup + others

Members in blue - attended March 2009 workshop

"Carlos Bertulani" <carlos_bertulani@tamu-commerce.edu>,
"David Bernholdt" <bernholdtde@ornl.gov>,
"Alex Brown" <brown@nscl.msu.edu>,
"Zaochun Gao" <gao1z@cmich.edu>,
"Gaute Hagen" <hageng@ornl.gov>,
"Mihai Horoi" <horoi@phy.cmich.edu>,
"Calvin Johnson" <cjohnson@sciences.sdsu.edu>,
"Plamen Krastev" <pkrastev@sciences.sdsu.edu>,
"Piotr Magierski" <magiersk@if.pu.edu.pl>,
"Pieter Maris" <pmaris@iastate.edu>,
"Petr Navratil" <navratil1@llnl.gov>,
"Esmond Ng" <egng@lbl.gov>,
"Erich Ormand" <ormand1@llnl.gov>,
"Sofia Quaglioni" <quaglioni1@llnl.gov>,
"Arturo Samana" <arturo_samana@tamu-commerce.edu>,
"Masha Sosonkina" < masha@scl.ameslab.gov>,
"Roman Sen'kov" <senkov@nscl.msu.edu>,
"Mario Stoitsov" <mario.stoitsov@gmail.com>,
"James Vary" <jvary@iastate.edu>,
"Chao Yang" <CYang@lbl.gov>

LCCI Meetings:

Pack Forest, August 2007
Pack Forest, June 2008
Oakland DNP, October 26, 2008
San Diego State University, March 12-14, 2009
Pack Forest, June 2009
Next meeting: DNP meeting in Hawaii

SciDAC/UNEDF LCCI Code/Environment Project

Vision Statement (Draft)

Researchers, from the level of advanced graduate students,
will have access to forefront research tools in theoretical physics
and related training material

Mission Statement (Draft)

To develop a flexible, user-friendly and sustainable code/environment
providing researchers, including students,
with convenient access to LCCI codes and archived results

LCCI Code/Environment Project

Goals (Draft)

- Increase the “discovery potential” of nuclear physics
- Enhance nation’s return on investment (ROI) in research and facilities
- Provide advanced training tools to help develop our S&T workforce
- Achieve a sustainable environment for high-value digital resources
- Increase nuclear physics’ outreach to other fields of science

Motivations (Draft)

- Help define/achieve SciDAC/UNEDF goals - years 4 & 5
- Help define/achieve critical research agenda under SciDAC III
- Help define/achieve next decade’s research agenda
  - e.g. DOE extreme scale white paper, next NP Long Range Plan

Framework for Actions (Draft)

- Self-identified network of SciDAC/UNEDF members + volunteers
- Meet ~2x/year
- Interact with UNEDF Council/Membership (obtain feedback & volunteers)

UNEDF SciDAC Collaboration

Universal Nuclear Energy Density Functional

LCCI deliverable will enable much more!

QCD
Theory of strong interactions

χEFT
Chiral Effective Field Theory

LCCI/UNEDF

Inter-Nucleon
NN, NNN Interactions
AV18, EFT, $V_{low\text{-}k}$

Theory of Light Nuclei
Spectroscopy and selected reactions
Verification: NCSM=GFMC=CC
Validation: nuclei with A<16

Big Bang
Nucleosynthesis
& Stellar Reactions

Density Functional Theory
improved functionals
remove computationally-imposed constraints
properties for all nuclei with A>16

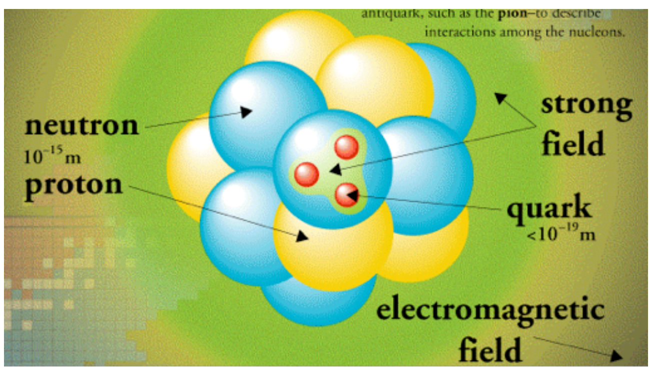

LCCI

Dynamic Extensions of DFT
LACM by GCM,TDDFT,QRPA
Level densites

Nuclei as
neutrino physics
laboratories

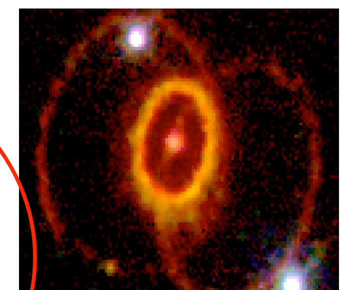

Low-energy Reactions
Hauser-Feshbach
Feshbach-Kerman-Koonin
Fission
mass and energy distributions

r,s processes
& Supernovae

www.unedf.org

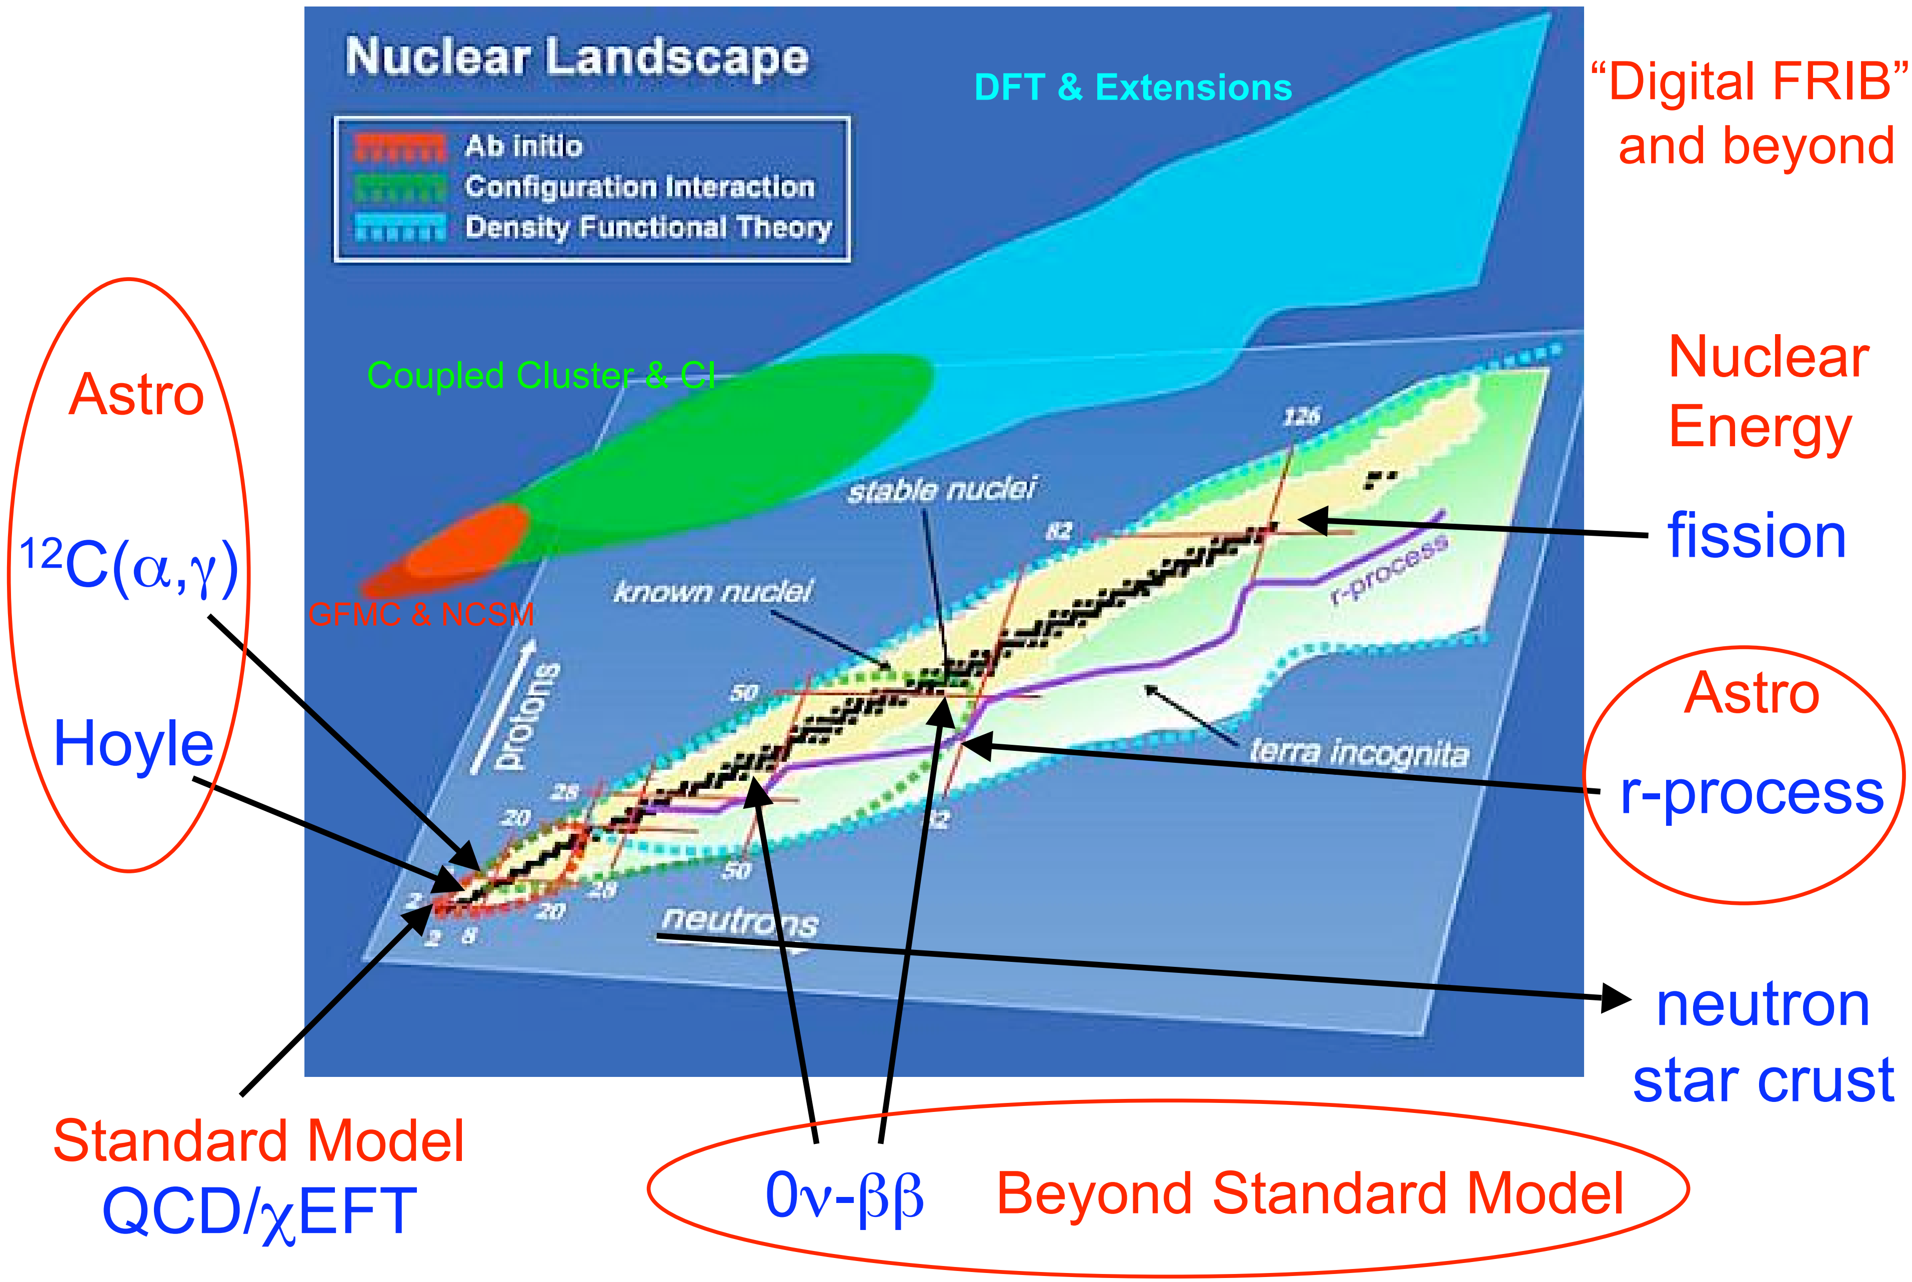

Nuclear Landscape
Ab initio
Configuration Interaction
Density Functional Theory
DFT & Extensions
Coupled Cluster & CI
GFMC & NCSM
stable nuclei
known nuclei
terra incognita
r-process
protons
neutrons
126
82
50
28
20
8
2
"Digital FRIB" and beyond
Nuclear Energy
fission
Astro
r-process
neutron star crust
Astro
$^{12}C(\alpha,\gamma)$
Hoyle
Standard Model
QCD/χEFT
0ν-$\beta\beta$
Beyond Standard Model

DOE Workshop on Forefront Questions in Nuclear Science
and the Role of High Performance Computing,
Gaithersburg, MD, January 26-28, 2009
Nuclear Structure and Nuclear Reactions

List of Priority Research Directions

- **Physics of extreme neutron-rich nuclei and matter**
- **Microscopic description of nuclear fission**
- **Nuclei as neutrino physics laboratories**
- **Reactions that made us - triple α process and $^{12}C(\alpha,\gamma)^{16}O$**

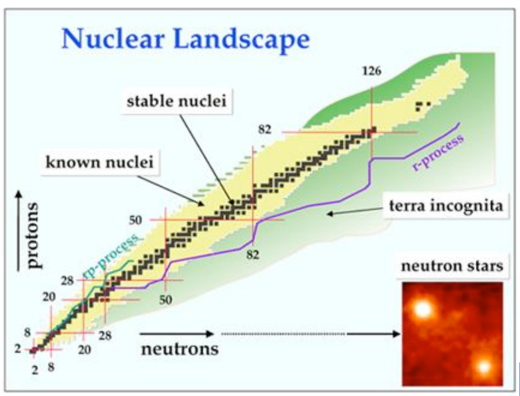

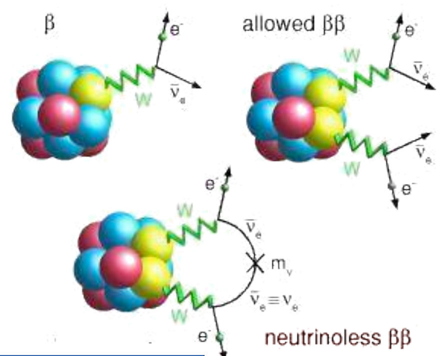

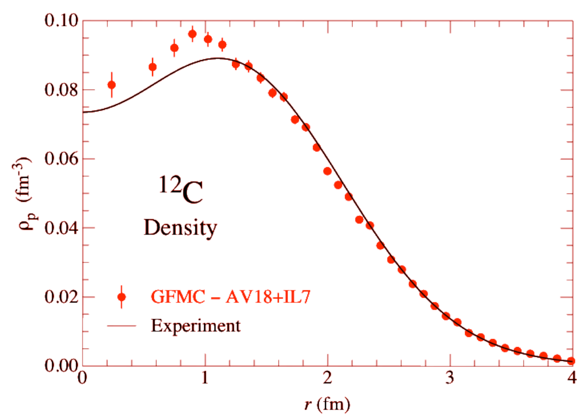

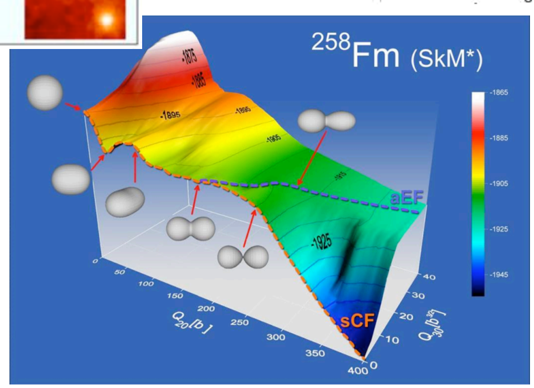

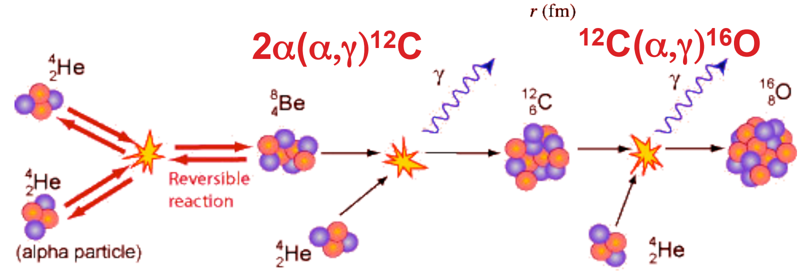

Physics of Extreme Neutron-Rich Nuclei and Matter

Scientific and computational challenges

- Understand structure of neutron-rich nuclei for r-process nucleosynthesis
- Solve structure and dynamics of neutron-star crust for cooling and gravity waves
- Implement load balanced, scalable algorithms, and include three-nucleon forces

Summary of research direction

- Perform ab initio calculation of closed-shell nuclei and their neighbors
- Predict shell structure and explore neutron-drip line
- Establish the nuclear Hamiltonian
- Develop algorithms for dynamics and transport in nucleonic matter

Expected Scientific and Computational Outcomes

- Determine limits of nuclear stability
- Determine static and transport properties of nucleonic matter
- Develop scalable algorithms for the quantum many-body problem

Potential impact on Nuclear Science

- Quantitative understanding of the origin of the elements from iron to uranium
- Drive experimental program at FRIB
- Determine properties of most dense nucleonic matter in the universe

DOE Workshop on Forefront Questions in Nuclear Science
and the Role of High Performance Computing,
Gaithersburg, MD, January 26-28, 2009
Nuclear Structure and Nuclear Reactions

Nuclei as Neutrino Physics Laboratories

Scientific and computational challenges

- Develop effective interactions and weak currents based upon fundamental theory and experiments.
- Diagonalize matrices of dimension $10^{12\text{-}13}$

Summary of research direction

- Develop extreme scale nuclear structure codes; estimate uncertainties with competing methods.
- Create techniques for ν-reactions relevant to oscillation experiments.

Expected Scientific and Computational Outcomes

- Dependence of $0\nu\ \beta\beta$-decay nuclear lifetimes on neutrino mass with theoretical uncertainty to 30-50%.
- ν-nucleus cross sections to 20%.
- Fault-tolerant, load-balanced highly scalable sparse eigensolver; load balancing for Monte Carlo simulations.

Potential impact on Nuclear Science

- Interpret experiments to explain the nature of the neutrinos and their masses.
- Calculate rates of nuclear reactions that drive stars and stellar explosions.

Reactions that Made Us - Triple α Process and $^{12}C(\alpha,\gamma)^{16}O$

Scientific and computational challenges

- First precise calculation of $2\alpha(\alpha, \gamma)^{12}C$ and $^{12}C(\alpha,\gamma)^{16}O$ rates for stellar burning
- Work done on one node must be split to a cluster of nodes
- Global sums

Summary of research direction

- Ab initio calculations of:
  - Hoyle state in ^{12}C
  - $^{8}Be(\alpha,\gamma)^{12}C$
  - States in ^{16}O
  - $^{12}C(\alpha,\gamma)^{16}O$

Expected Scientific and Computational Outcomes

- Accurate rates of two key reactions
- Ab initio tool for low-energy light-ion reactions and weakly-bound nuclei
- Library for distributing shared-memory work to subsets of massively parallel machine

Potential impact on Nuclear Science

- Provides essential input to modeling of stellar evolution and element production
- Provides a firm basis for extrapolating future experimental results
- Guide and interpret light exotic nuclei studies at FRIB
- We will know we truly exist!

Nuclear Physics Requires Exa-scale Computation

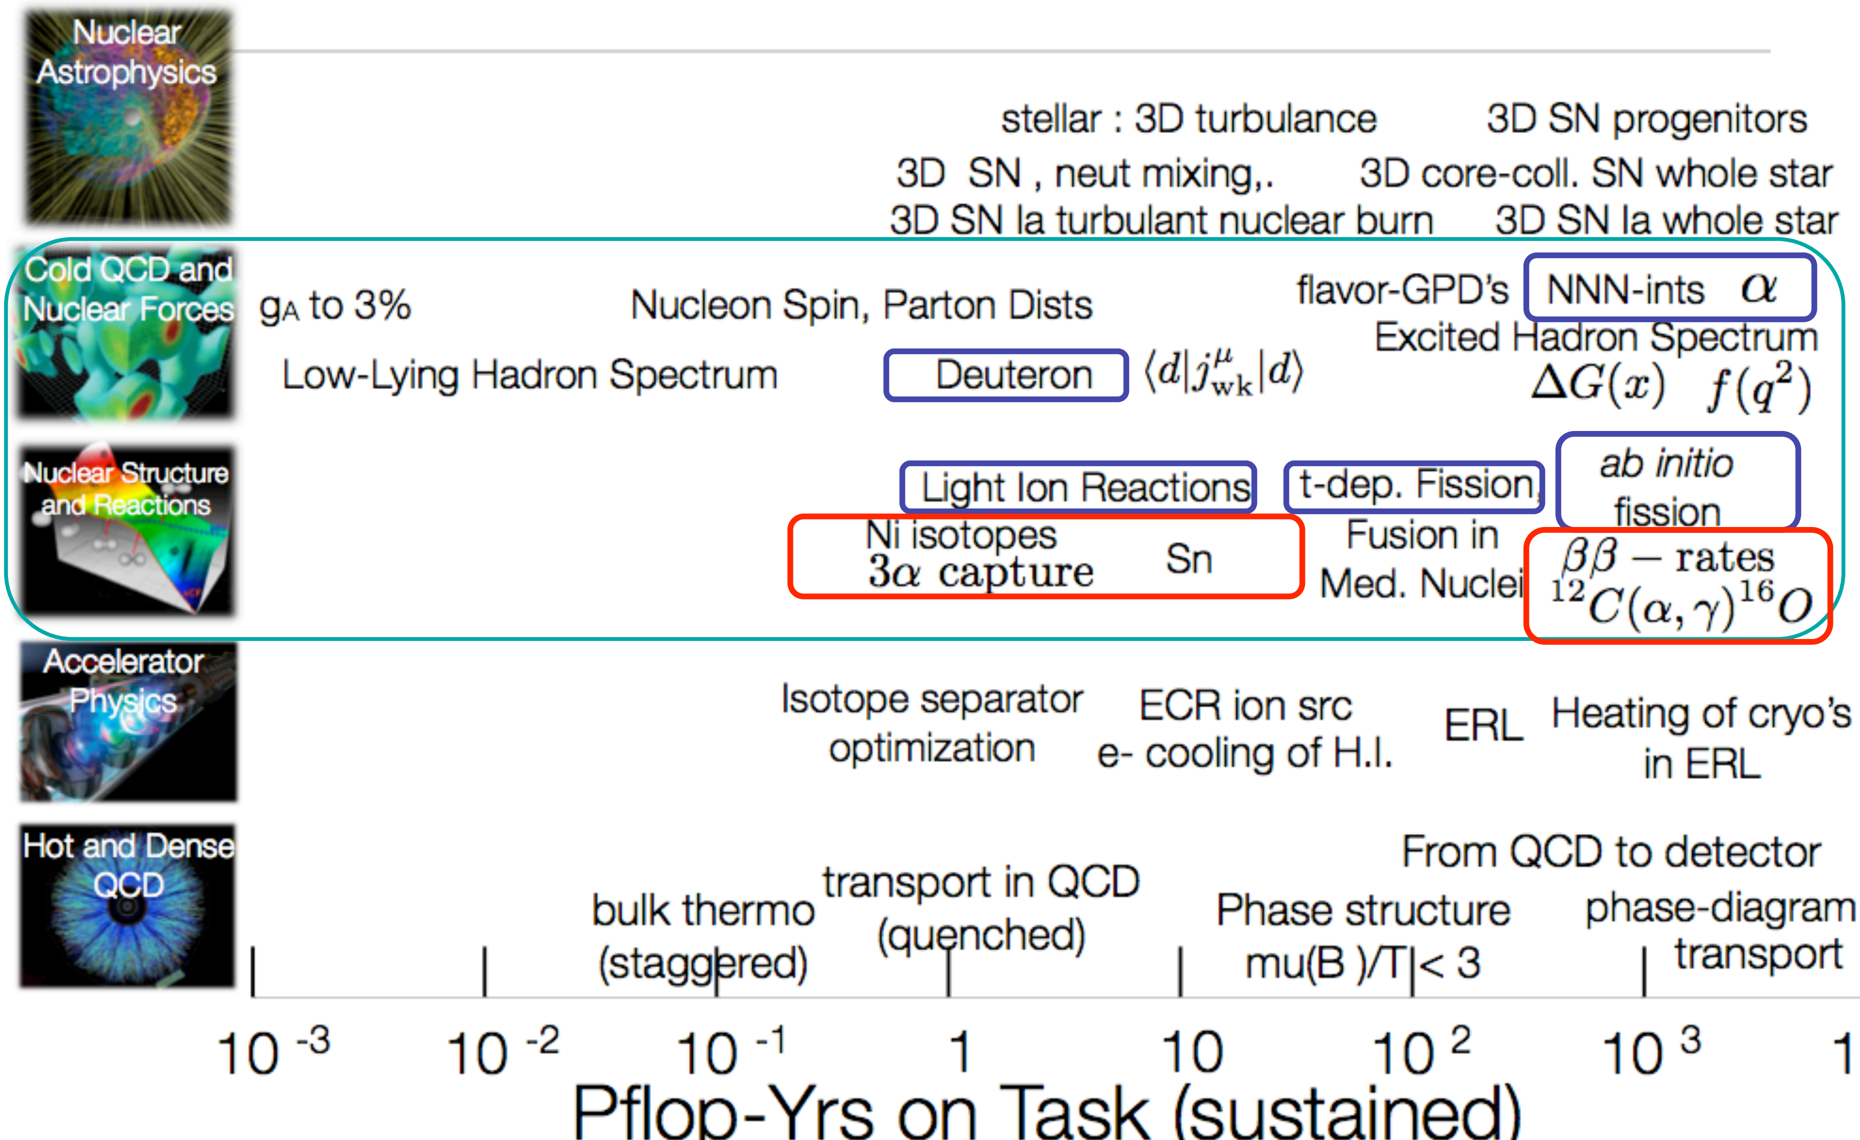

SciDAC/UNEDF related

Augmented "reach" of LCCI effort

Exa-scale computing will unify Nuclear Physics

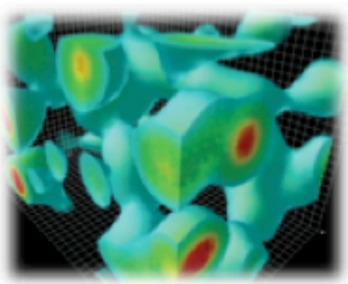

Cold QCD and Nuclear Forces

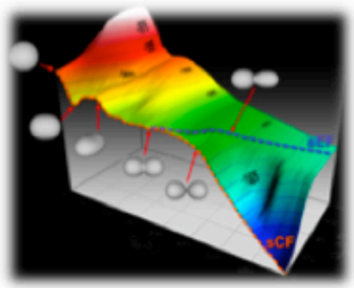

Nuclear Structure and Reactions

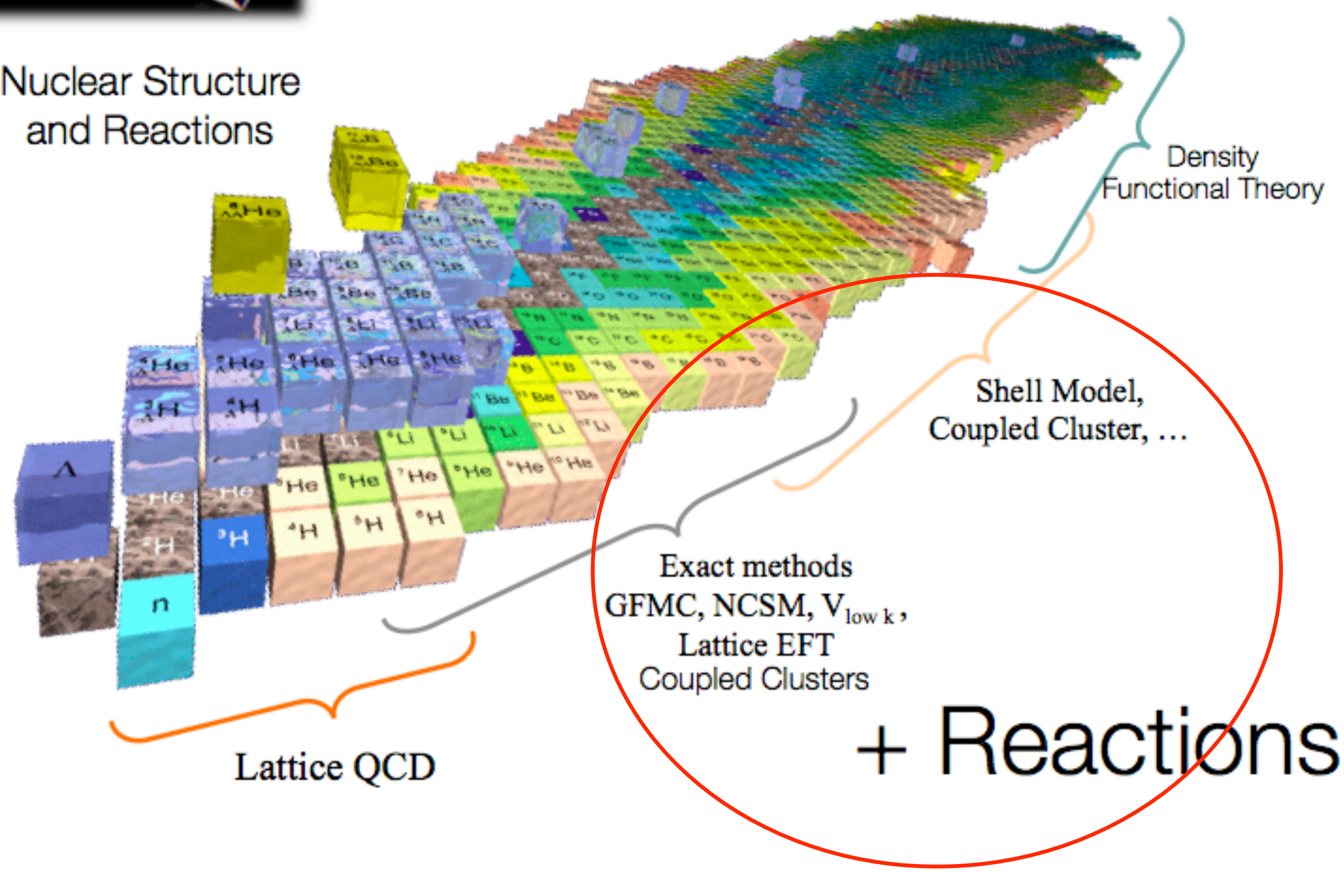

Human Resources

- To fully utilize Exa-scale at 2017 we need to grow expertise in the NP community
    - faster than Moores Law
    - not business as usual

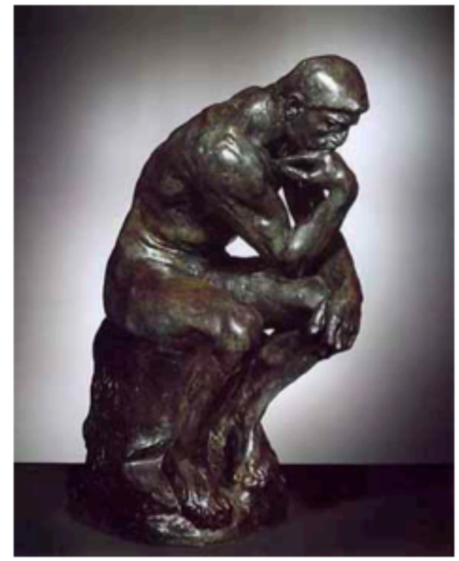

- The standard interdisciplinary hiring problems exist
    - challenge at Labs and Universities
    - new training models , start today for 2017?

- Broad collaborations
    - Graduate students and postdocs hired into collaboration
        - naive scaling from RHIC and UNEDF programs = SIGNIFICANT enhancement in person-power (+10+10 per project ?)

    - Organization in the Nuclear Physics community

Workshop Goals

- “Look under the hood” - technical exchanges
- Advance the plans for LCCI code/environment
- Volunteer for prototyping specific sub-elements
- Develop our strategy for moving forward

Some questions to begin addressing

- Do we wish to achieve an “official” UNEDF status?
  Ans: We will report our status and plans and seek feedback
- Do we release our 5 working documents to UNEDF via website?
  Ans: After a polished draft is achieved, we will place it on the UNEDF website behind password protection and invite feedback

Workshop Program

Thursday March 12
9 - 9:45 am: James Vary: discussion of overall goals
9:45 - noon: Vary/Maris: matrix storage algorithms and MFD
** lunch **
1:30 - 2:30: Navratil: TRDENS density matrix code and importance sampling
2:30 - 4:30: Ormand/Johnson: on-the-fly algorithms and REDSTICK
4:30 - 5:30 : Analysis and discussions

~ 6 pm: dinner at Mexican restaurant in Old Town

Friday March 13
9 - 9:45 Discussions and planning
9:45 - 11 am: Bernholdt (by telecon) - Common Component Architecture (CCA)
11 - noon: Horoi: Hash tables and other topics
** lunch **
1:30 -3 pm : Ng and Yang, lead: discussion of computational issues
3 pm break
3:30-4:30: James Vary colloquium for SDSU Computational Sciences Research Center

~ Social with SDSU Computational Science Department followed by dinner at Johnson-Perdue home

Saturday March 14
9am - noon: ALL: initial specifications for code ; review “5-reports”, outline development of code (e.g. basis, Hamiltonian, etc) and assignments

LCCI White Paper - Working Drafts

Introduction (Ormand)
Physics drivers (Johnson)
Physics capabilities (Maris)
Module/environment (Ormand, Johnson)
Technical features (Horoi)
Work flow of advanced CI code (Johnson)

Example of a multi-code environment

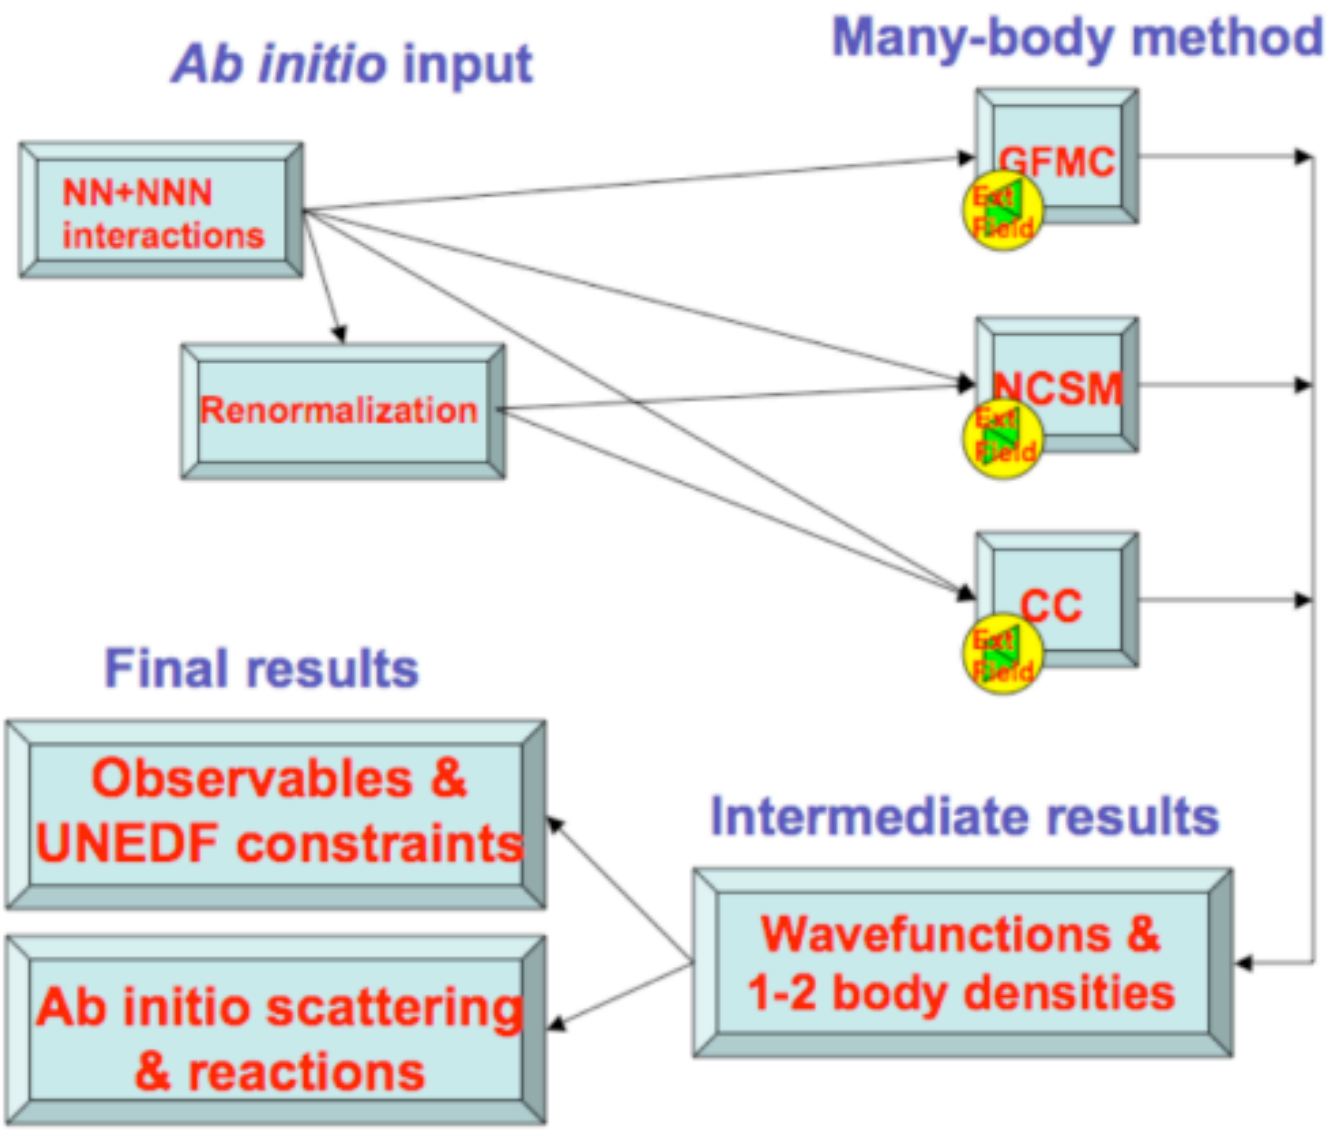

LCCI Environment - Overview

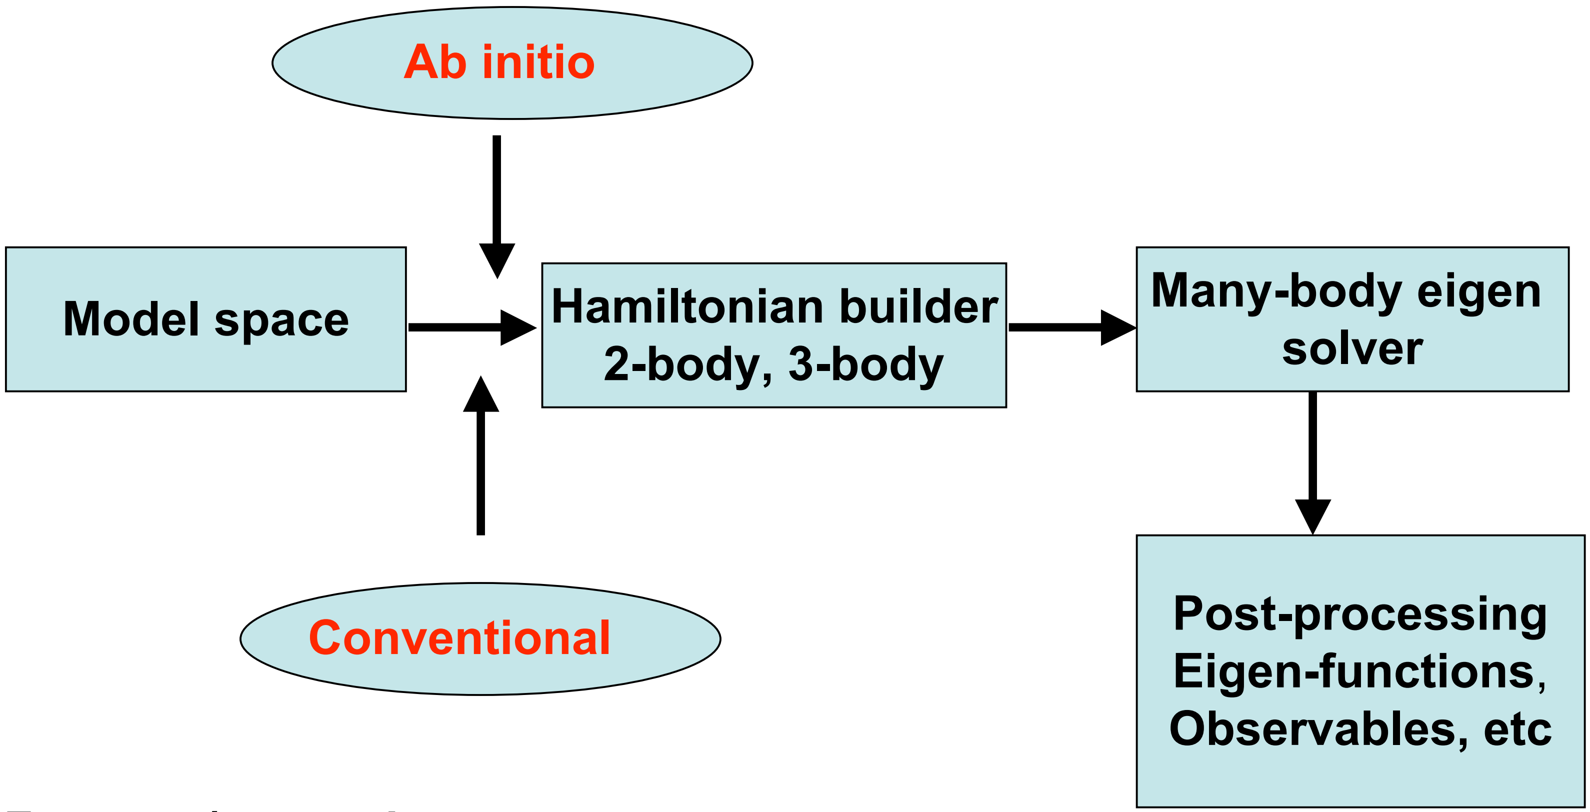

Features/properties:

Scalable 1-10^6 cores

Goal: LCCI code should be capable of solving, at least, the identified consensus structure problems from DOE “extreme scale” white paper

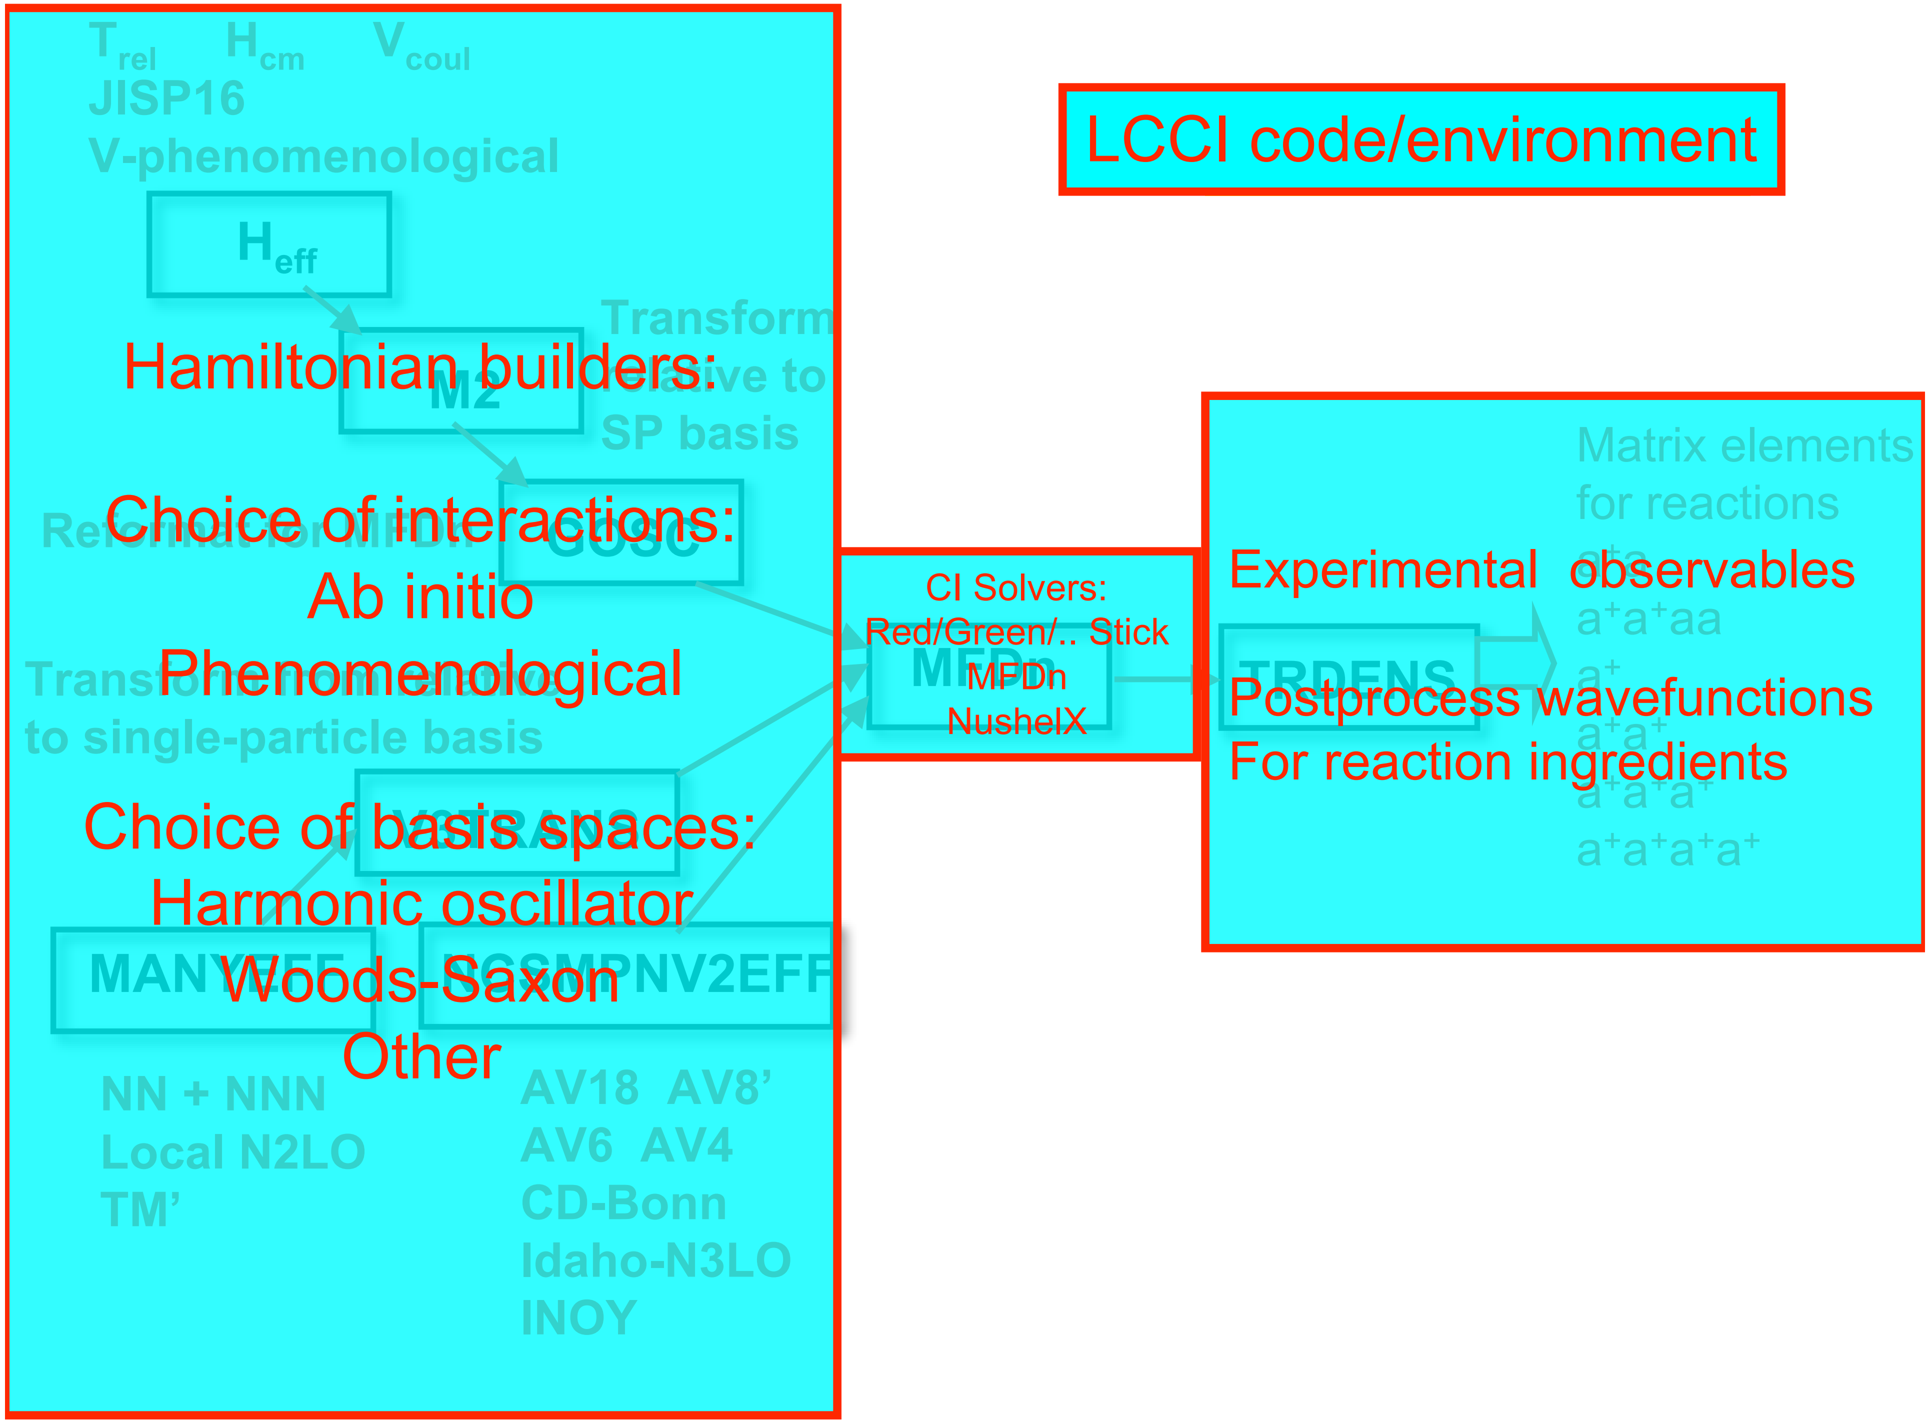

LCCI code/environment
Hamiltonian builders:
Choice of interactions:
Ab initio
Phenomenological
Choice of basis spaces:
Harmonic oscillator
Woods-Saxon
Other
CI Solvers:
Red/Green/.. Stick
MFDn
NushelX
Experimental observables
Postprocess wavefunctions
For reaction ingredients
T rel H cm V coul
JISP16
V-phenomenological
H eff
NN + NNN
Local N2LO
TM’
AV18 AV8’
AV6 AV4
CD-Bonn
Idaho-N3LO
INOY
to single-particle basis
SP basis
MANYEFF
Matrix elements
for reactions

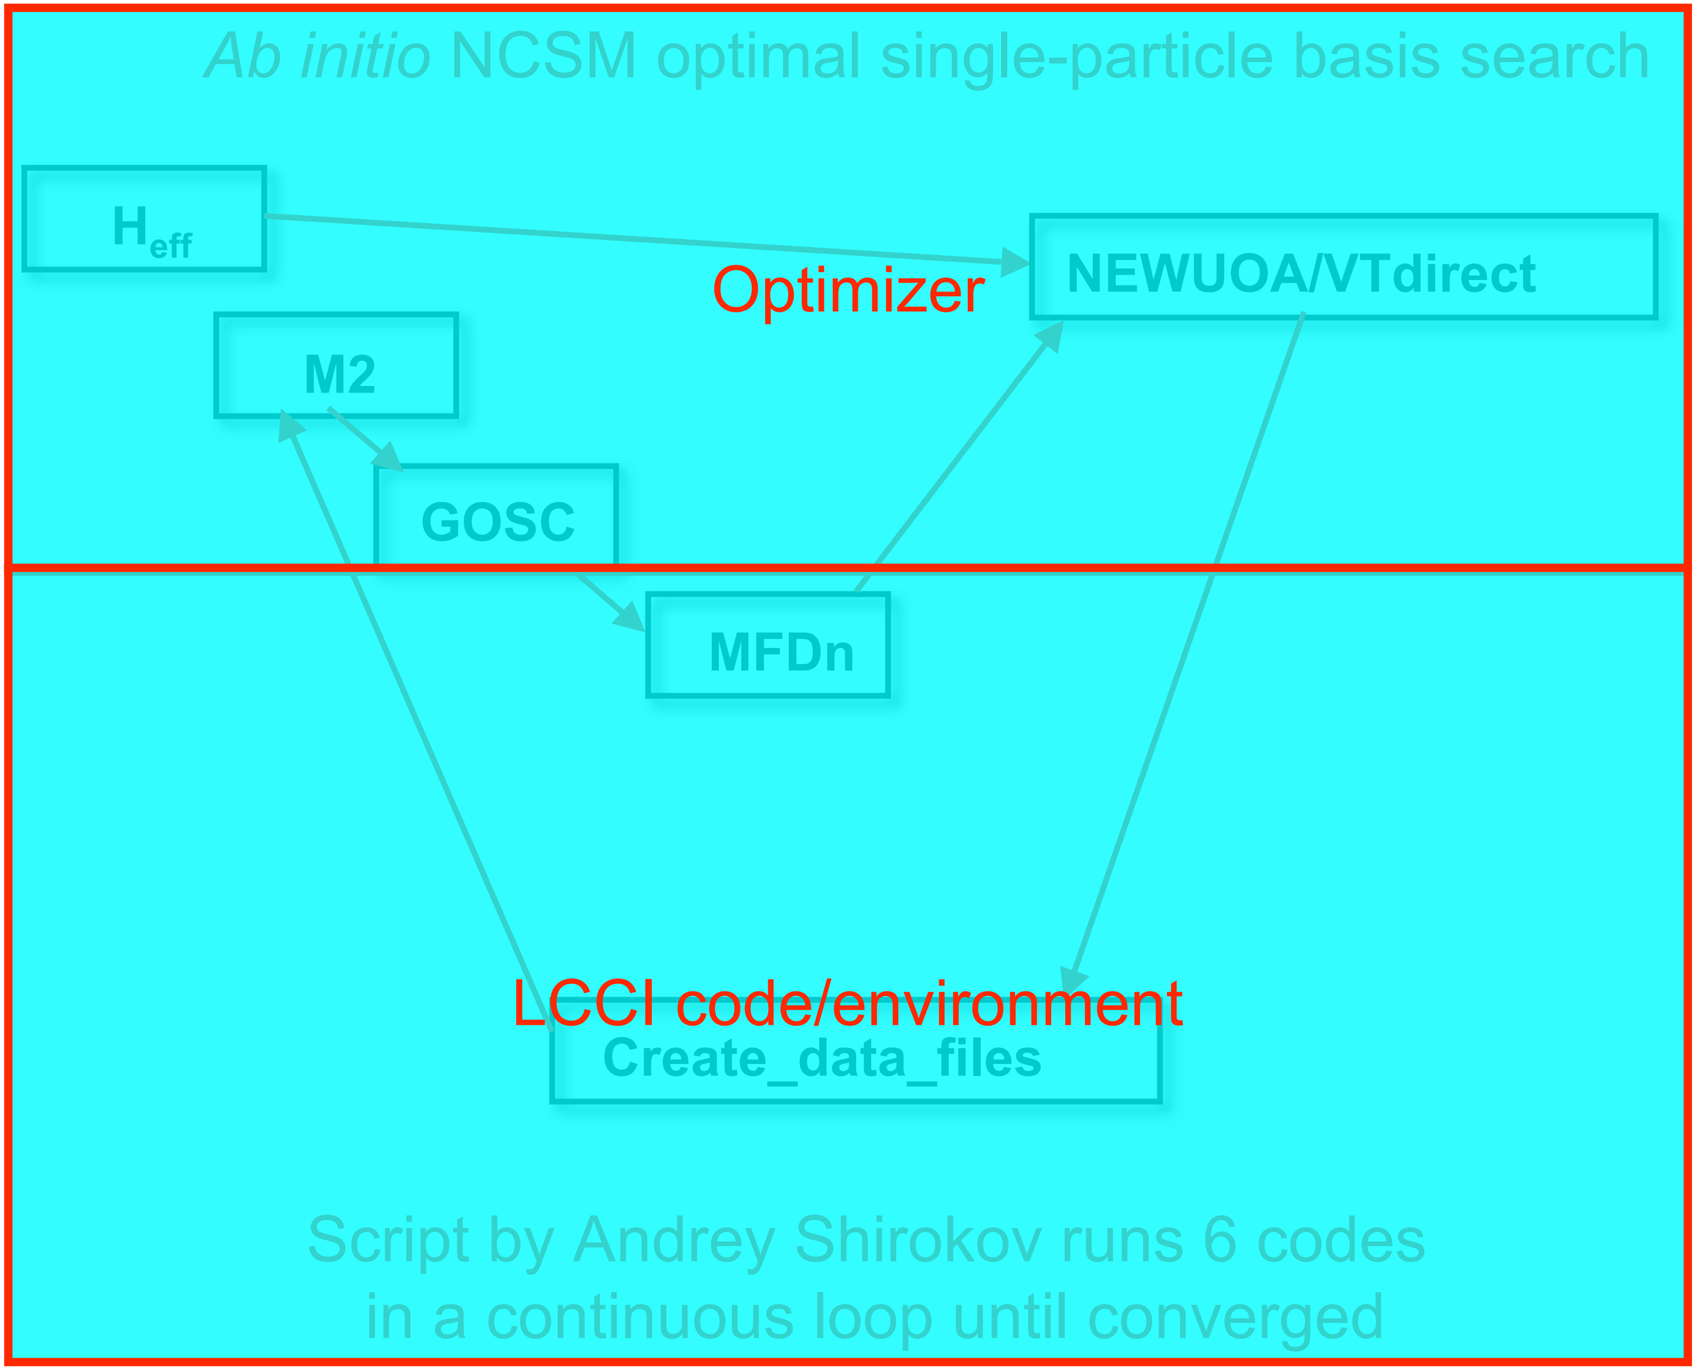

Ab initio NCSM optimal single-particle basis search
H_{eff}
M2
GOSC
MFDn
Optimizer
NEWUOA/VTdirect
LCCI code/environment
Create_data_files
Script by Andrey Shirokov runs 6 codes
in a continuous loop until converged

What is the user-friendly environment that we envision?

Examples to provoke discussion:

- Primitive example - nuclear physics server at ISU
- Advanced example - ECCE (PNL)

Cooperation in the Digital Age

- Archive data & programs for distribution
- Upload once for everyone
- Keep track of updates
- Leave notes
- Host of “Nuclear Physics Calculator”

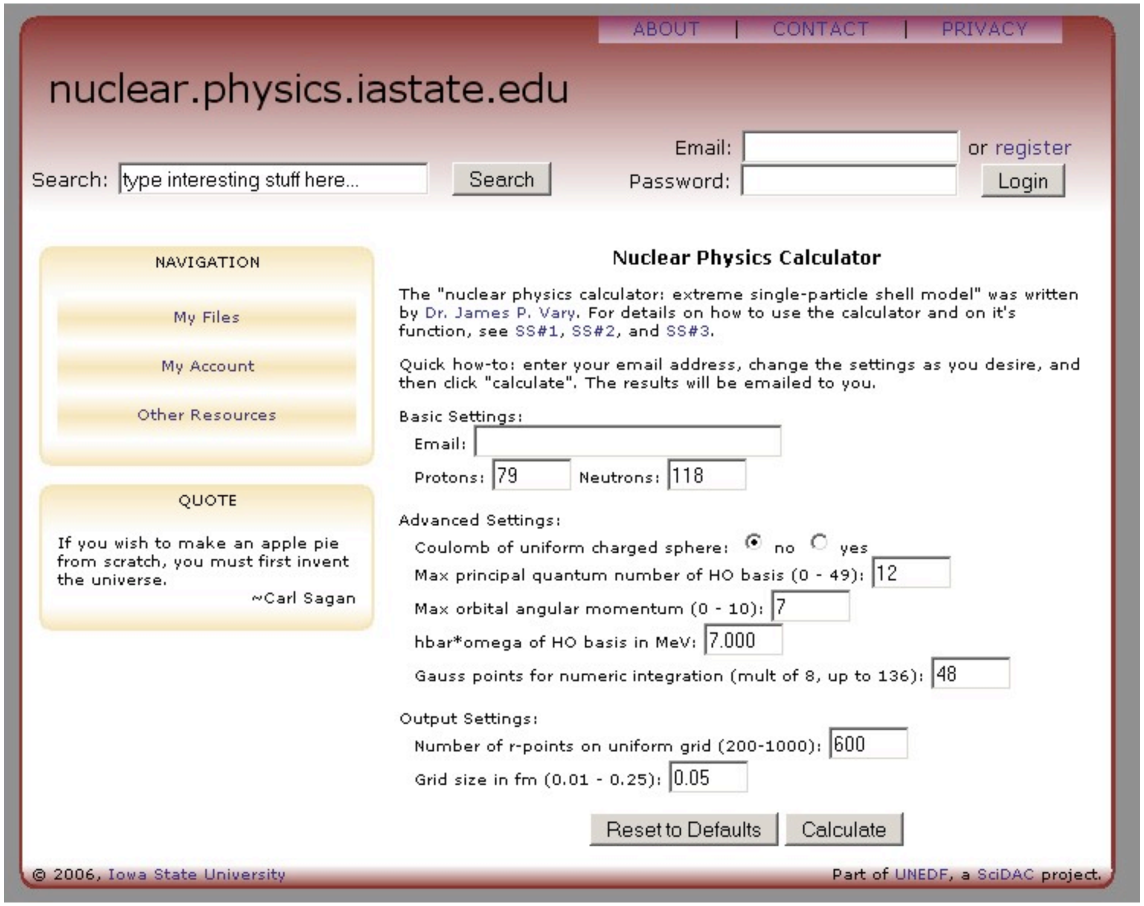

Nuclear Physics Calculator: *ab initio* NCSM demonstration project

nuclear.physics.iastate.edu

- Select the NCSM application
- Enter your email address
- Select the number of neutrons and protons
- Select the N_{max}
- Select the oscillator energy, $\hbar\Omega$

Results file with the JISP16 interaction will be emailed to you in a few minutes

Note that this is a demonstration project with limited single-processor calculations and will evolve as funding permits

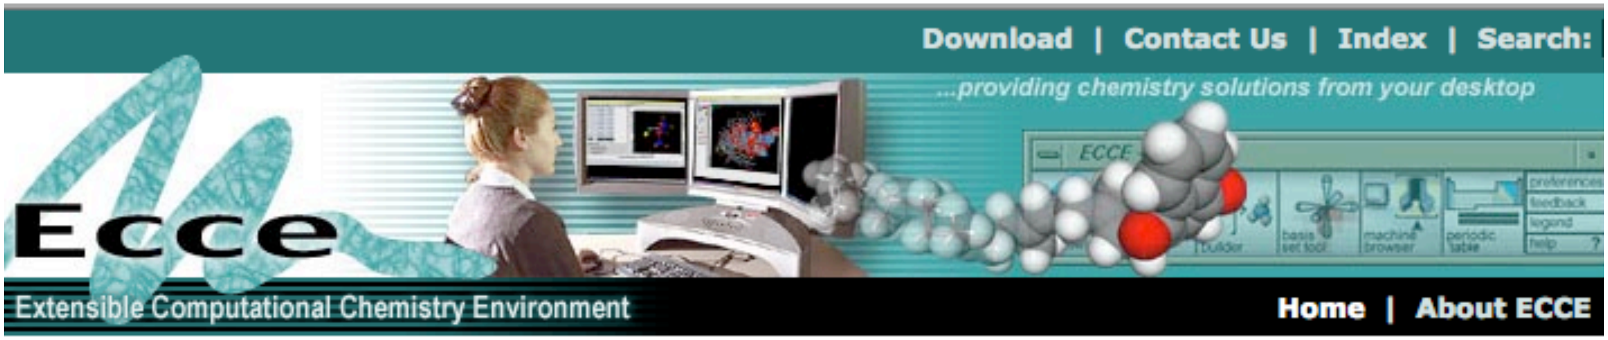

The Extensible Computational Chemistry Environment (ECCE, pronounced "etch-ā") provides a sophisticated graphical user interface, scientific visualization tools, and the underlying data management framework enabling scientists to efficiently set up calculations and store, retrieve, and analyze the rapidly growing volumes of data produced by computational chemistry studies.

General Features

- Support for building molecular models.
- Graphical user interface to a broad range of electronic structure theory types. Supported codes currently include NWChem, GAMESS-UK, Gaussian 03™, Gaussian 98™, and Amica. Other codes are registered based on user requirements.
- Graphical user interface for basis set selection.
- Remote submission of calculations to UNIX and Linux workstations, Linux clusters, and supercomputers. Supported queue management systems include PBS™, LSF™, NQE/NQS™, LoadLeveler™ and Maui Scheduler.
- Three-dimensional visualization and graphical display of molecular data properties while jobs are running and after completion. Molecular orbitals and vibrational frequencies are among the properties displayed.
- Support for importing results from NWChem, Gaussian 94™, Gaussian 98™, and Gaussian 03™ calculations run outside of the ECCE environment.
- Extensive web-based help.

The ECCE application software currently runs on Linux workstations and is written in C++ using the wxWidgets user interface toolkit and OpenGL graphics. Ongoing development will extend ECCE to predict reaction rate constants as well as build and run calculations on periodic systems.

What are the standards/conventions/protocols for

- review of codes before “adoption”
- documentation and standard test cases
- program interfaces
- code updates/extensions
- data archiving/indexing/release
- read/write access privileges
- reviewing/prioritizing requested improvements

Do we need to invent all these ourselves or have useful models been developed by others (e.g. CCA, large experimental collaborations, other SciDACs,…)?

AM/CS Issues in Nuclear CI Calculations

- Eigenvalue calculations
  - –Compute leading edge of the eigenvalues of extremely large sparse matrices
  - –Optimal scalable parallel eigenvalue solvers
- More compact representation of operators/functions and data
  - –E.g., 3-body input files and many-body basis states
- Sparse matrix-vector multiplications.
  - –Efficient factorization-based MV multiplications.
  - –Structured matrices (dense blocks).
- Alternative formulation or solution methods
  - –E.g., J-projections.

AM/CS Issues in Nuclear CI Calculations

- Massive global nonlinear optimization with nonlinear constraints
  - –E.g., finding optimal basis states and effective interactions.
- Combinatorial mathematics
  - –Whitehead's approach
  - –Optimal partitioning scheme for minimizing searches
  - –Bit manipulations
- Compute traces?

AM/CS Issues in Nuclear CI Calculations

- Load balancing
  - ADLB?
- Programming models
  - Interaction between multiple address spaces
  - Programming languages (Fortran, C, C++)
  - Libraries, Global Arrays, UPC, MPI+x, Multithreading, …
  - Parallel (asynchronous) I/O
- Fault tolerance - very important on very large no. of cores
  - Application level vs system level, latency level, communications, …
- File I/O (checkpoint/restart, local/global)
- Debugging tools and performance analysis tools
  - Performance analysis extremely useful and important

AM/CS Issues in Nuclear CI Calculations

- Validation/verification
  - –Code comparison (in terms of results).
  - –Comparisons with experimental results
- Code benchmarking
  - –Performance comparison among CI codes
- Productivity (to improve workflow)
  - –CCA???
  - –Programming environment, coding, …
  - –User interface environment (e.g., ECCE for quantum chemistry)
- Massive data management
  - –Input/output, archiving
  - –Visualization, data analysis

Overview of Workshop Outcomes

- Shared Technologies
  - "Peek under the hood" of each major code
- Defined CS/AM issues in Nuclear CI Calculations
- Explored algorithms that may be adopted/modified to benefit the LCCI code
  - Hashing methods, Whitehead combinatorics, PARPACK, Locally optimized preconditioned conjugate gradient
- Investigated potential shared components/algorithms
  - REDSTICK <-> MFDn, NuShellx -> REDSTICK,
  - interface CI codes with TRDENS
- Investigated potential shared code development environments - CCA, ECCE
- Discussed pros/cons of developer's environment - NERSC vs UNEDF server vs ?
- Planned next steps
  - Report writing - unify "5 reports" into single report with intro/overview
  - Planned separate page on the UNEDF website (UNEDF-only)
    - with workshop talks if requested
  - Protocols/standards/conventions
  - Proactive connections to other components of UNEDF

San Diego Workshop Action agenda

- Develop draft workshop report and circulate to participants, edit, submit to UNEDF website (JPV)
  Done - year 3 accomplishment

- Edit draft “5 reports”, polish and submit to UNEDF website (behind PW protection) (all)
  In process - year 3 accomplishment

- Draft intro to “5 reports” (WEO with input from PM & CJ)
  In process - year 3 accomplishment

- Investigate cross-use of Redstick’s - MFDn’s basis generation/storage methods (WEO, CJ, JPV, PM)
  Done - year 3 accomplishment

LCCI - year 4 plans

Developments leading to a first generation UNEDF LCCI environment

Prioritization - what does the user need first?
How to provide for those needs?
More ?'s

Specified at today's LCCI meeting - all were welcome and resulted in the list on the next page

Working group lists planned:
LCCI group (exists)
Users group(s) (by signup in project space at NERSC)

Develop/First generation code/data environment available to UNEDF researchers via NSERC project space
(Groups: MSU/CMU/ISU/LLNL/SDSU/LBNL/AL)

NushellX+script/data/testruns installed/tested (Group MSU/CMU)

MFDn + script/data/testruns installed/tested
(Group ISU/LBNL/AL)

RedStick + script/data/testruns (Group LLNL/SDSU)

Obtain/generate/test elements of the environment
(MSU/CMU/ISU/LLNL/SDSU)

Develop timeline at Hawaii meeting
(MSU/CMU/ISU/LLNL/SDSU/LBNL/AL)
