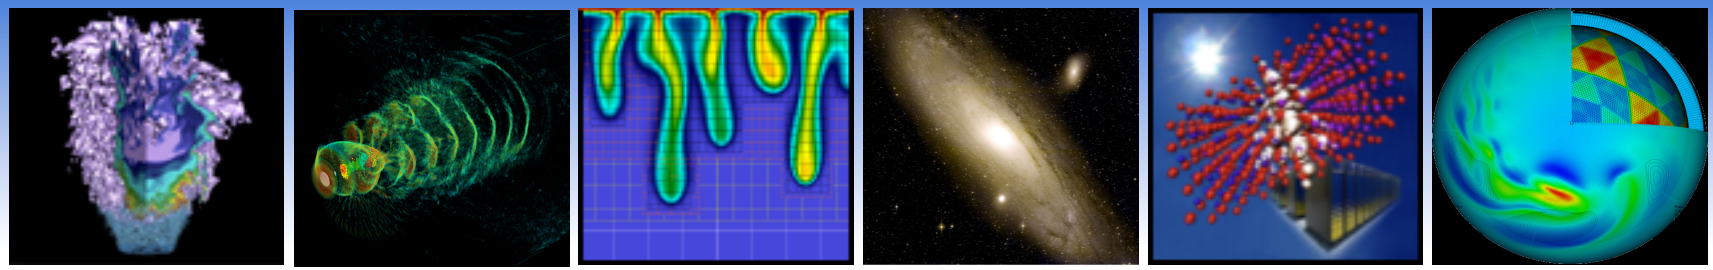

Workshop Goals & Process

Large Scale Computing and Storage Requirements for Nuclear Physics Research

NP / ASCR / NERSC Workshop

May 26-27, 2011

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC National Energy Research Scientific Computing Center

BERKELEY LAB

Logistics: Schedule

- Agenda on workshop web page
  - https://www.nersc.gov/science/requirements-workshops/nuclear-physics/agenda
  - Need your presentation slides
- Mid-morning / afternoon break, lunch
- Self-organization for dinner

Friday Morning Schedule

- Some time available for case studies that don't finish on Thursday
- Richard & Harvey summary
- Further discussion case studies: how many (if needed); content
  - Initial table entries
- Richard & Harvey available to help
- Lunch available?

Workshop Content

- 6 “science areas,” one workshop
  - Science-focused but cross-science discussion; Explore areas of common need
  - Low energy nuclear physics
  - Lattice QCD
  - Nuclear astrophysics
  - Heavy ion experiments
  - Nuclear data
  - Nuclear accelerators

Why is NERSC Collecting Computational Requirements?

- NERSC is science driven.
- Your input helps create the science-based justification for
  - acquiring the resources and
  - implementing the services that you need to reach your research goals
- Help NERSC make informed decisions for technology and services
  - guide procurements, staffing, and to improve the effectiveness of NERSC services.

What Input is Sought?

- Includes hardware, software, support, data, storage, analysis, work flow
  - Time frame: 2014
- Different from ERCAP:
  - Longer term focus
  - Not what you think you can get, but what you need

Logistics: Case Studies

- Narrative describing science & NERSC reqmts
- Audience is NERSC, DOE program managers
- Minimum set to capture NP mission and unique NERSC requirements
  - Initial set suggested by Ted Barnes
  - Some science areas may have more than one method in use
  - Different science areas may use the same method
  -

Examples of Information Sought

- Type of simulation, number of, reason for #, algorithms, solver
- Parallelism: method, weak or strong scaling, implementation, concurrency, limits
- Key physical parameters and their limits:
  - spatial resolution, # of atoms/energy levels, integration range, …
- Representative code
- Key science result metrics and goals

Examples of Information Sought

- Typical science process (workflow)
- Data: amount stored / transferred for input, results, and fault mitigation
- Special needs for data intensive projects
  - Grids, gateways, workflows, provenance, `
- Special query regarding multicore/manycore

- How all of this is
  - Driven by the science
  - Likely to change and why

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Final Thoughts

- Requirements characterization process is not complicated.
- Mutually beneficial.

Scaling Science

Inspired by **P. Kent**, "*Computational Challenges in Nanoscience: an* ab initio *Perspective*", Peta08 workshop, Hawaii (2008) and **Jonathan Carter** (NERSC).

Length, Spatial extent, #Atoms, *Weak scaling*

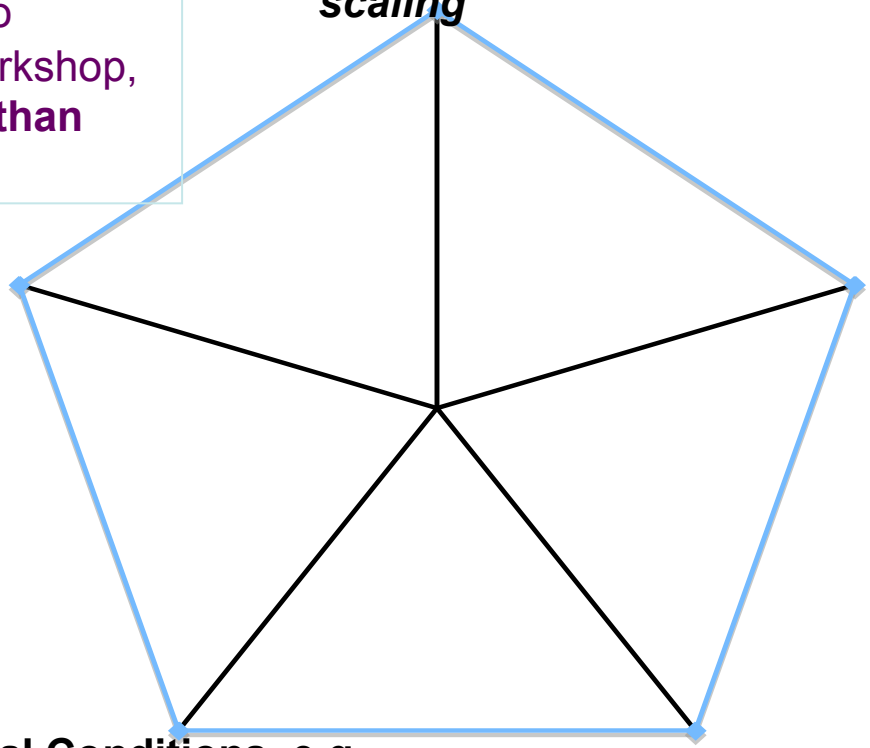

Time scale Optimizations, *Strong scaling*

Convergence, systematic errors due to cutoffs, etc.

Initial Conditions, e.g. molecule, boundaries, *Ensembles*

Simulation method, e.g. DFT, QMC or HF/ SCF; LES or DNS