Present and Future Computing Needs in Math Software Research & Development

Esmond G. Ng

Lawrence Berkeley National Laboratory

January 6, 2011

Project Overview

- Case Study Requirement Worksheet focused on an allocation titled "High Performance Sparse Matrix Algorithms"
  - this allocation supports sparse matrix R&D related to 4 SciDAC projects.
    - TOPS ("Towards Optimal Petascale Simulations")
      - including collaborations with EFRC's through SciDAC-e funding
    - ComPASS ("Community Petascale Project for Accelerator Science and Simulation")
    - UNEDF ("Building a Universal Nuclear Energy Density Functional")
    - CEMM ("Center for Extended Magnetohydrodynamic Modeling")
  - **not** all activities in TOPS, ComPASS, UNEDF, and Fusion
    - just activities at LBNL

- Main goal is to develop highly efficient and scalable sparse matrix problems, with an emphasis on highly indefinite and ill-conditioned linear systems, as well as eigenvalue problems

Current HPC ~~Methods~~ Research

- Sparse linear systems
  - direct methods
    - based on triangular factorizations … developed (efficient, scalable) SuperLU_DIST
      - memory bound …
    - On-going improvements to SuperLU_DIST
  - preconditioned iterative methods
    - focused on factorization-based preconditioners
    - versions of SuperLU (and eventually SuperLU_DIST) that perform incomplete factorization
    - PDSLin – a new parallel linear solver that is based on domain decomposition and employs the Schur complement methods
      - hybrid approach that combines direct and iterative methods
- Sparse eignevalue calculations
  - tailoring/improving Lanczos algorithms (including those in PARPACK) for different applications
  - develop new/alternative eigenvalue algorithms

Current HPC Requirements

- Codes are currently based on MPI; R&D entirely on NERSC systems
- Need parallel graph partitioning tools in linear solvers
  - ParMetis, PTScotch
  - ordering, domain decomposition, …
- Also need PETSc, LAPACK, SLAPACK, and BLAS

- HPC usage is modest
  - mostly for algorithmic development and testing
  - occasionally need to perform scaling studies and performance evaluations, in which case we may need larger core counts (<10,000 so far)

- HPC requirements can be very different from the application scientists' perspective
  - depend on the problems they are solving

Limitations, Obstacles, and Bottlenecks

- Large memory requirements
  - some of the matrix problems from applications are extremely large
  - sparse direct solvers are memory bound
  - data locality is important
- Can be communication bound
  - e.g., sparse triangular solutions
- Load balancing
  - related to problem partitioning
  - there are other issues
- May involve quite a bit of integer computation
- Data read – limited to matrix input [probably unnecessary in "production" runs]
  - depend on matrix size

HPC Usage and Methods for the Next 3-5 Years

- Still expect to be relatively modest for algorithmic research and development
  - core counts expected to be higher (>10,000)

- R&D direction
  - codes are currently MPI based, but we are beginning to look at mixed programming model
  - significant algorithmic changes are expected to accommodate multi-/many-cores architectures
  - dynamic load balancing may be needed

HPC Usage and Methods for the Next 3-5 Years

- Fast turnaround very desirable for testing and evaluation purposes
- Desirable to have some processing nodes to have much more memory than other nodes
- Tools for performance profiling, tuning, and optimization are essential on future architectures
- Strategy for heterogeneous architectures (including those with accelerators):
  - HELP ☹

New Science with New Resources

- ❑ We work on enabling technology …
- ❑ More memory, larger core count in the system (and larger allocation) will allow our sparse matrix solvers to handle much larger problems
- ❑ This in turn will enable scientific applications to carry out much larger scale modeling and simulation
  - and there are such applications

Solving CEBAF BBU Using Shape Uncertainty Quantification Method

SciDAC Success as a Collaboration between Accelerator Simulation, Computational Science and Experiment – Beam Breakup (BBU) instabilities at well below the designed beam current were observed in the CEBAF12 GeV upgrade of the Jefferson Lab (TJNAF) in which Higher Order Modes (HOM) with exceptionally *high* quality factor (Q) were measured. Using the shape uncertainty quantification tool developed under SciDAC, the problem was found to be a deformation of the cavity shape due to fabrication errors. This discovery was achieved as a team effort between SLAC, TOPS, and JLab which underscores the importance of the SciDAC multidisciplinary approach in tacking challenging applications.

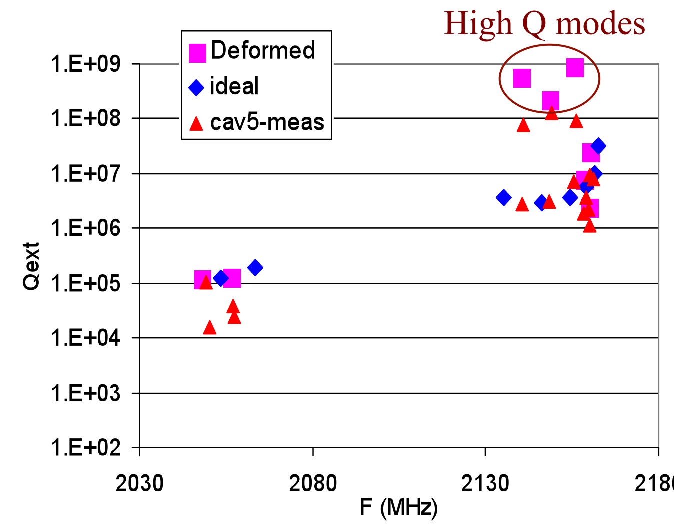

Method of Solution - Using the measured cavity parameters as inputs, the deformed cavity shape was recovered by solving the ***inverse*** problem through an optimization method. The calculations showed that the cavity was 8 mm shorter than designed, which was subsequently confirmed by measurements. The result explains why the troublesome modes have high Qs because in the deformed cavity, the fields shift away from the HOM coupler where they can be damped. This shows that quality control in cavity fabrication can play an important role in accelerator performance. .

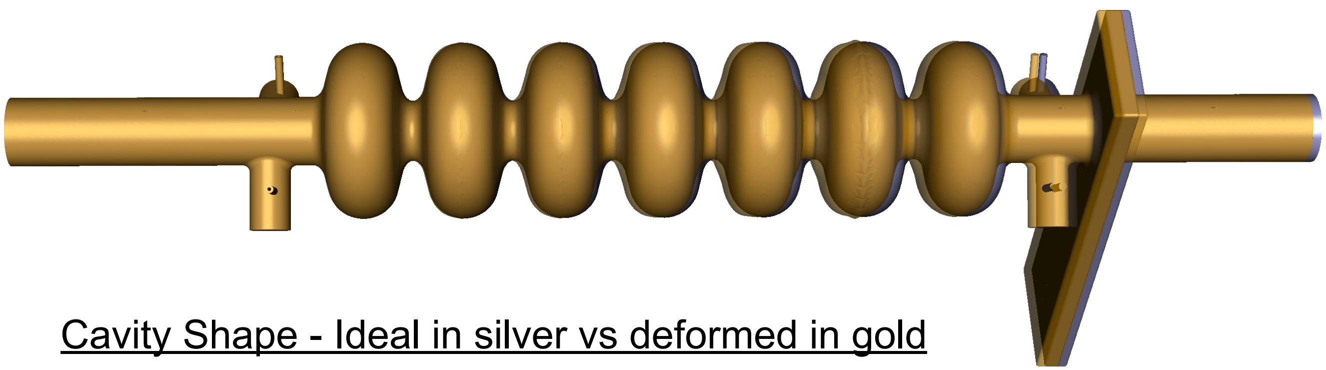

Cavity Shape - Ideal in silver vs deformed in gold

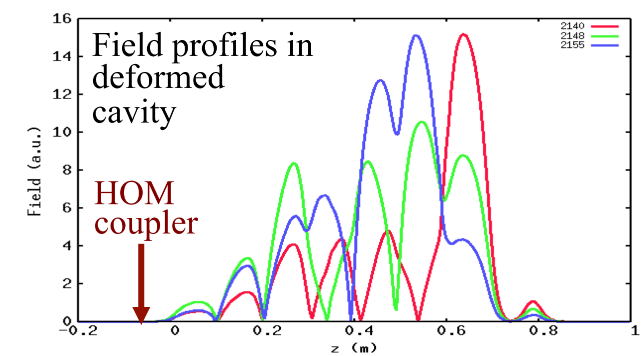

Proton-Dripping Fluorine-14

First principles quantum solution for yet-to-be-measured unstable nucleus ^{14}F

- ❖ Apply *ab initio* microscopic nuclear theory's predictive power to major test case
- ❖ Robust predictions important for improved energy sources
- ❖ Providing important guidance for DOE-supported experiments
- ❖ Comparison with new experiment will improve theory of strong interactions
- ❖ Dimension of matrix solved for 14 lowest states ~ 2×10^9
- ❖ Solution takes ~ 2.5 hours on 30,000 cores (Cray XT4 Jaguar at ORNL)

Predictions:

Binding energy: 72 ± 4 MeV indicating that Fluorine-14 will emit (drip) one proton to produce more stable Oxygen-13.

Predicted spectrum (Extrapolation B) for Fluorine-14 which is nearly identical with predicted spectrum of its "mirror" nucleus Boron-14. Experimental data exist only for Boron-14 (far right column).

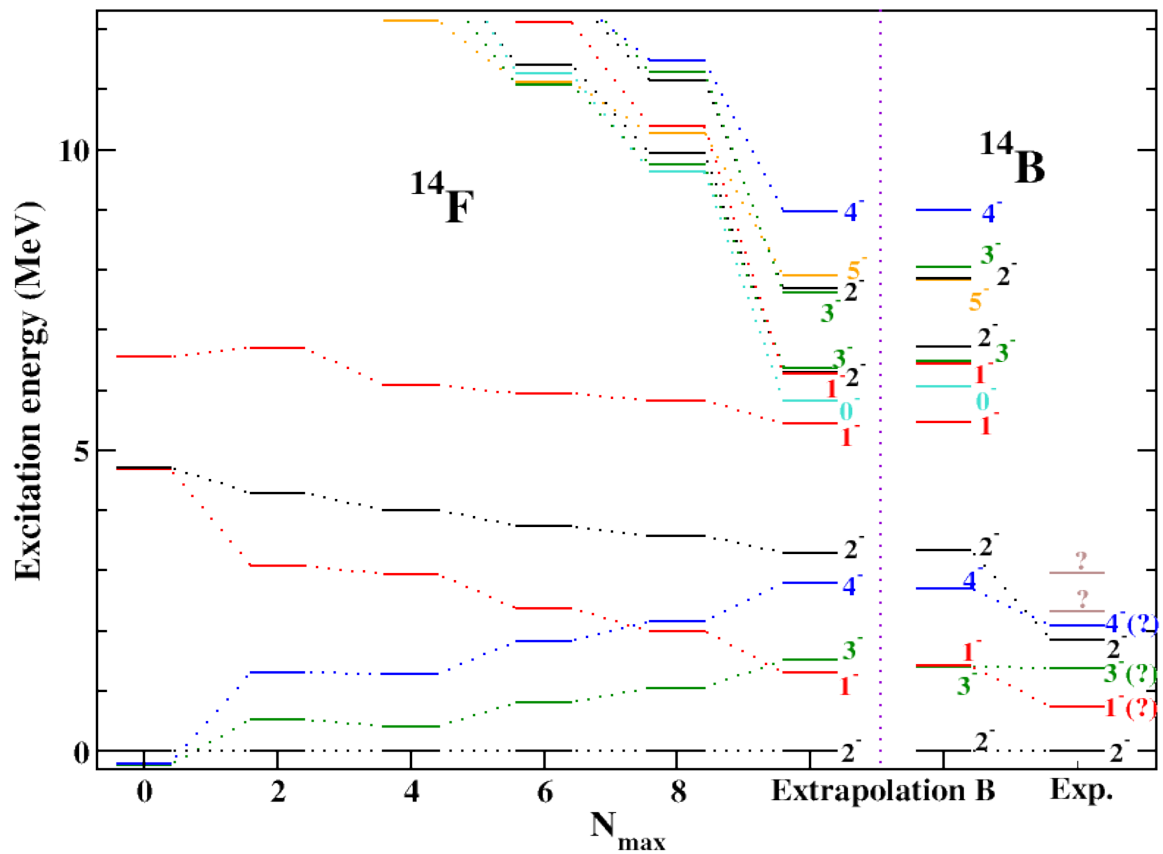

New Science with New Resources

- Challenges in large-scale nuclear structure calculations

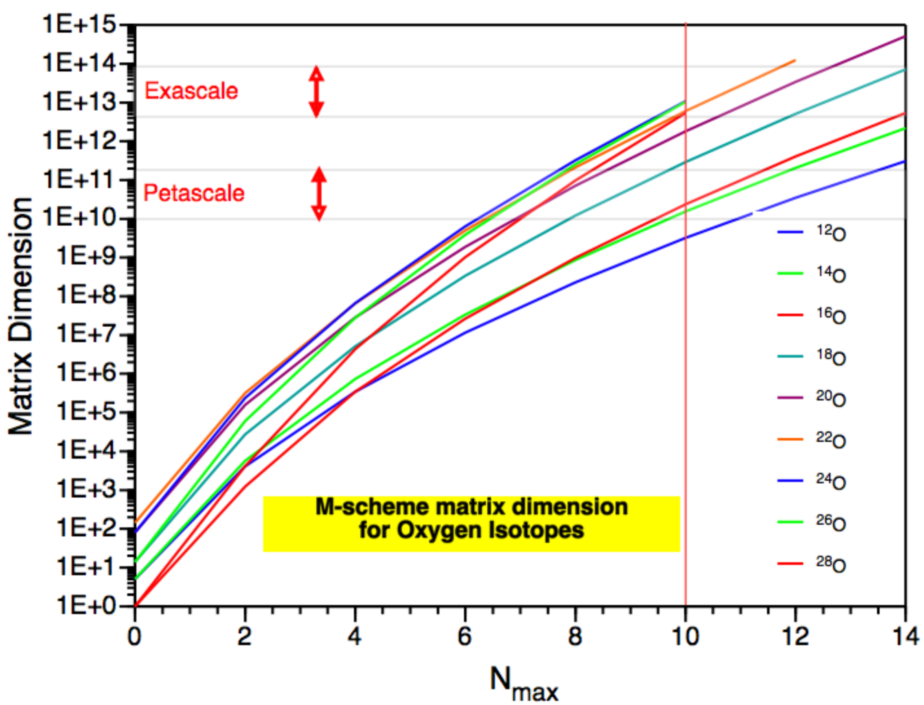

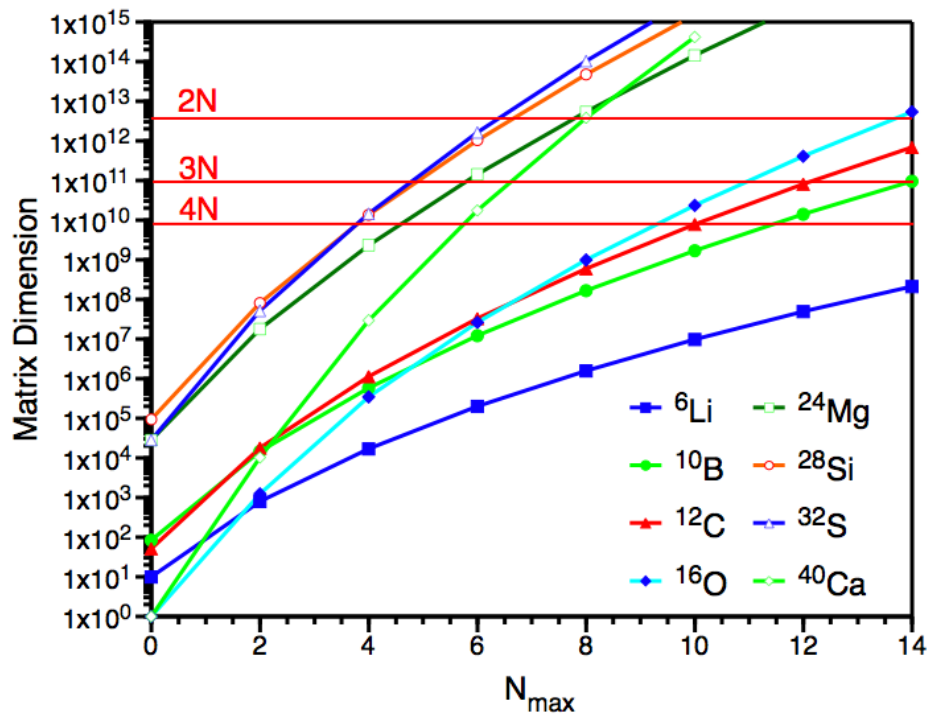

New Science with New Resources

- Challenges in large-scale nuclear structure calculations