Present and Future Computing Requirements for Macroscopic ITER Dynamics

Carl Sovinec,[1] Stephen Jardin,[2] and Linda Sugiyama[3]

[1]University of Wisconsin-Madison

[2]Princeton Plasma Physics Laboratory

[3]Massachusetts Institute of Technology

NERSC BER Requirements for 2017
September 11-12, 2012
Rockville, MD

1. Project Description

Stephen Jardin/PPPL (CEMM director)
Alan Glasser/Univ. WA
Eric Held/USU
Val Izzo/UCSD
Scott Kruger/Tech-X Corp.
Scott Parker/Univ. CO at Boulder
Jesus Ramos/MIT
Dalton Schnack/Univ. WI-Madison
Mark Shephard/RPI
Carl Sovinec/Univ. WI-Madison
Henry Strauss/HRS Fusion
Linda Sugiyama/MIT

- Macroscopic (magnetohydrodynamic-like) activity is problematic for magnetic confinement of plasma:
  - Performance-limiting magnetic islands,
  - High transient heat flux from edge-localized modes,
  - Discharge termination and material damage associated with disruptive activity.
- Scientific progress in the ITER experiment (~$20B, 1st plasma in 2020) will depend on control of macroscopic dynamics.
- Simulation-based modeling is critical for understanding, avoidance, and control.

1. Project Description (continued)

A central objective of CEMM is validated predictive simulation of macroscopic phenomena in tokamaks.

- Present-day studies tend to be phenomenological.
  - Modeling provides insight into how dynamics occur and which effects are important.
  - Practical considerations limit the separation of time- and space-scales, relative to currently operating experiments.
- By 2017 advances in hardware and software will enable quantitative validation at realistic parameters.
  - The dimensionless Lundquist number (S=resistive diffusion time/Alfvén-wave propagation time) is a reasonable measure, and scale-separation ~ S^{α}, $1/3 \leq \alpha \leq 1/2$.
  - Achieving realistic values of $S = 10^7$-10^8 is feasible.
  - Future modeling will also have greater emphasis on two-fluid and kinetic effects and on data analysis.

2. Computational Strategies

- Our computations solve initial-value problems in 3 spatial dimensions, evolving plasma number density, flow velocity, magnetic field, and temperatures.
  - The MHD and two-fluid models are nonlinear PDEs.
  - Problems of interest are stiff.
- The primary codes of CEMM are M3D, M3D-*C1*, and NIMROD.
  - To address temporal stiffness, the codes use implicit methods.
  - To address anisotropy, geometry, and constraints, high-order finite-element, spectral-element, and spectral methods are used.
- Our codes are characterized by these algorithms:
  - Parallel domain decomposition (mostly MPI) in 3 spatial dimensions,
  - Numerical integration of projections onto finite and spectral bases,
  - Generation of large sparse ill-conditioned matrices, and
  - Krylov-space iterative solution with parallel direct solution of sparse diagonal blocks for preconditioning.

2. Computational Strategies (continued)

- Characteristic problem sizes have 1000s-10,000s of high-order poloidal (2D) basis functions with 16-64 nodes per element and 16-64+ toroidal planes or toroidal Fourier components.
- Computations require 1000s-100,000 time-steps.
- Our biggest computational challenge is parallel solution of the sparse ill-conditioned matrices at each time-step.
  - It is also the primary limitation of parallel scaling.
- Memory limitation: preconditioning via direct methods over blocks does not yield linear scaling with problem size.
- By 2017, we expect to
  - Extend daily and full-scale parallelization by factors > 10 each,
  - Incorporate modeling for plasma kinetics, and
  - Perform finite-time Lyapunov exponent (FTLE) analysis concurrently with simulations.

3. Current HPC Usage

- CEMM currently uses Hopper, Carver, and now Edison
- Hours used in 2012:
  - NERSC - ~4M (M3D), ~2M (NIMROD, divided among 4 repositories)
  - Local clusters - ~500,000
- Typical parallel concurrency and run time, number of runs per year
  - Computation size varies from 100 to 3000 cores.
  - Many (5-10) smaller computations are needed for each large full-resolution computation.
  - Wallclock hours for a single computation is 24-96.
  - The number of computations per year is ~200.
- Data read/written per run
  - Data read per run is minimal
  - Data written per run is 1-100 GB.
  - Checkpoint file size is ~0.1 GB

3. Current HPC Usage (continued)

- Memory used per core is currently as large as 1.5 GB.
  - Aggregate memory is as large as 1 to 5 TB.
- Necessary software, services or infrastructure
  - PETSc, SuperLU_DIST, HDF5, ADIOS, VisIt, Fortran-90, and OpenMP
  - Increasing support for 3D visualization is important.
    - AVS-Express is better than VisIt for at least some applications.
- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored
  - HPSS total use ~25 TB
  - Global File System ~1.5 TB
  - On-Line file storage is approximately hourly for running jobs

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Compute hours needed (in units of Hopper hours)
    - Increasing the S-values by 10-100 from present-day computation requires approximately a factor of 4 increase in spatial resolution in two of the spatial dimensions and a factor of 2 in the third. Scaling from present-day, CEMM will require approximately 200 M hours.
- Changes to parallel concurrency, run time, number of runs per year
    - With improved scaling, concurrency will increase by approximately 10 for both small and large computations. Smaller computations will be ~1000 cores, and full-resolution computations will be ~100,000 cores. Run time and number of runs will remain at current levels.
    - FTLE analysis may need longer run times to characterize structures that form from turbulence.
- Changes to data read/written
    - Data and storage will increase linearly with resolution (~32).
- Changes to memory needed per core
    - Increases to 4 GB per core will help avoid memory limitations.

5. Strategies for New Architectures

- Significant performance increases with new architectures will require porting of parallel algebraic system solvers.
  - PETSc (collection of methods)
  - Stand-alone sparse direct solvers (SuperLU_DIST)
  - Native solvers with relatively simple preconditioning may become more effective, but large iteration count is problematic for orthogonalization (GMRES)
- We have investigated MPI/OpenMP hybrid parallelization, finding improvement on computations that are limited by memory with MPI alone.
- To date we have learned about case studies with GPU programming.
- We are planning to pay attention to new developments.
- Data analysis (FTLE) should be well suited for GPU computation to the extent that a vector field is fully available to a GPU.
- To be successful on many-core systems, we will need help.

6. Summary

- Improvements in NERSC computing hardware, software and services will facilitate simulations of macroscopic plasma dynamics at realistic parameters for model-validation.
- While the user community expects NERSC to adapt new architectures, flexibility and balance are important for applications requiring multi-faceted algorithms.
- Access to 32X current NERSC allocation will allow computation of internal plasma dynamics at realistic S-values of 10^7-10^8. It will also enable greater use of two-fluid and kinetic modeling.
- To date, our modeling has not been one of the more data-intensive applications, but storage and processing at the projected level will require more consideration.
- Provisions for long run times will facilitate analysis of turbulence and exploratory-scale computations.