NIMROD: Computational Aspects and Challenges

Ping Zhu

University of Wisconsin-Madison
with major contribution from C. R. Sovinec (UW-Madison)

CScADS Summer Workshop
Snowbird, UT
July 19-22, 2010

NIMROD: Non-ideal Magnetohydrodynamics with Rotation - Open Discussion

- NIMROD is an extended magnetohydrodynamics (MHD) code [Sovinec *et al.* 2004].
- NIMROD is a team project
  - Started in 1995-1996, involves team members from multiple institutions.
  - https://nimrodteam.org
- NIMROD is part of the DOE SciDAC centers CEMM (Center for Extended MHD Modeling) and SWIM (Simulation of Wave Interactions with MHD)
  - The other major MHD code in CEMM is M3D
  - Solve similar problems with somewhat different equations and schemes.

NIMROD code simulates macroscopic high temperature plasma dynamics

- Laboratory plasmas:
  - Magnetic fusion reactors: Tokamak, RFP, etc.
  - Laboratory astrophysics experiments: Reconnection, dynamo, etc.
- Space and astrophysical plasma
  - Magnetosphere and substorm
  - Accretion disk and jets

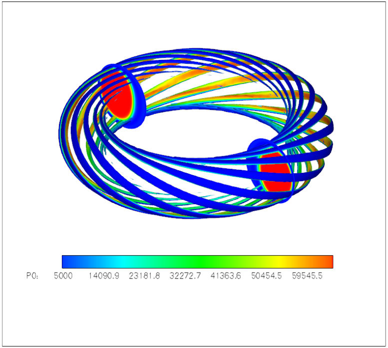

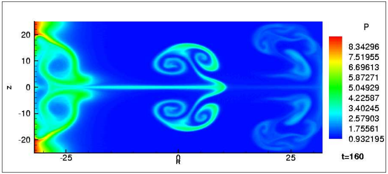

NIMROD code solves the extended MHD equations

- Fluid part:

$$\frac{d\rho}{dt} = -\rho\nabla\cdot\mathbf{u} + D\nabla^2\rho \quad (1)$$

$$\rho\frac{d\mathbf{u}}{dt} = -\nabla p + \mathbf{J}\times\mathbf{B} - \nabla\cdot\boldsymbol{\pi} \quad (2)$$

$$\frac{n}{\gamma-1}\frac{dT}{dt} = -\frac{p}{2}\nabla\cdot\mathbf{u} - \boldsymbol{\pi} : \nabla\mathbf{u} - \nabla\cdot\mathbf{q} + Q \quad (3)$$

$$\frac{\partial\mathbf{B}}{\partial t} = -\nabla\times\mathbf{E} \quad (4)$$

$$\mu_0\mathbf{J} = \nabla\times\mathbf{B} \quad (5)$$

$$\mathbf{E} = -\mathbf{u}\times\mathbf{B} + \eta\mathbf{J} + \frac{\lambda}{ne}(\mathbf{J}\times\mathbf{B} - \nabla p_e) \quad (6)$$

- Kinetic part: can couple to kinetic/particle code through moment closures.

Physics challenges for MHD simulations

- General macroscopic processes
  - Global geometry
  - Nonlinearities
- High temperature magnetized plasma specific
  - Stiffness:
    - Multiple temporal and spatial scales:

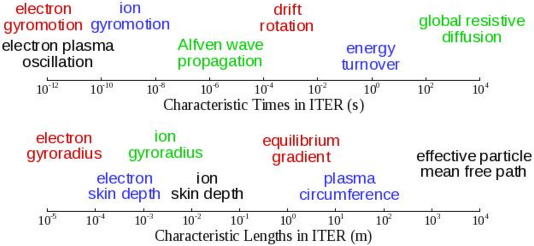

    - Coupling between fluid and kinetic scales
  - Anisotropy:
    - Extreme anisotropy imposed by magnetic field: $\kappa_{\parallel}/\kappa_{\perp} \sim 10^8$
  - Magnetic divergence constraint $\nabla \cdot \mathbf{B} = 0$

NIMROD uses mixed discretization schemes

- Spatial
  - High-order elements represent 2D poloidal domain.
    - Uniform collocation Lagrangian polynomials
    - Spectral elements with Gauss-Lobatto Legendre nodes
  - Finite Fourier series represents the periodic direction.
  - 3D matrices with inherently dominant diagonal blocks
- Temporal
  - Semi-implicit operator with predictor-corrector advance
  - Fully implicit operator with leap-frog advance
  - Fully implicit operator with time-centered advance
  - Allows large time step advance without numerical dissipation or instability

NIMROD solves large sparse matrices each step

- Linear problem: 2D matrix
  - Conjugate gradient or GMRES iterative solvers
  - Direct solver SuperLU [Li and Demmel, 2003] is used as preconditioner
- Nonlinear problem: 3D matrix
  - Matrix-free conjugate gradient (symm)
  - Matrix-free Krylov space GMRES (non-symm)
  - Fourier component block based preconditioning
    - Diagonal blocks: SuperLU
    - Limited off-diagonal blocks: Jacobi or Gauss-Seidel iteration

The main computational challenge is to find scalable solver for large ill-conditioned sparse matrix

- System stiffness leads to large ill-conditioned sparse matrix
- Numerically, the solution has been to use 2D and 3D preconditioning
  - SuperLU
  - Jacobi or Gauss-Seidel
- The challenge is to enable these solvers to be scalable in peta-scale computations.
- Efficient parallel computing schemes are also part of the solution.

NIMROD uses distributed memory paradigm and MPI communication

- Domain decompositions
  - Grid-block decomposition uses point-to-point communication
  - Fourier-Layer and configuaration-space decomposition
    - Domain swap between forward/inverse FFT
    - Collective communication
- Parallel communication during preconditioning
  - SuperLU_Dist
  - Jacobi/Gauss-Seidel: point-to-point communication

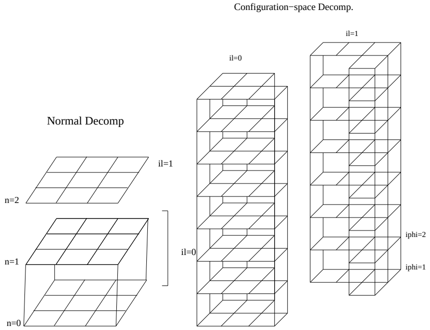

Parallel scaling of NIMROD code through 10K processor cores has been achieved [Sovinec 2009]

- Parallelization optimization
  - Overlap asynchronous communication for block Gauss-Seidel iteration with on-processor computation
  - Reorder loop and data allows fewer, larger collective communications
- Weak scaling through 10K processors on Franklin Cray XT-4 achieved

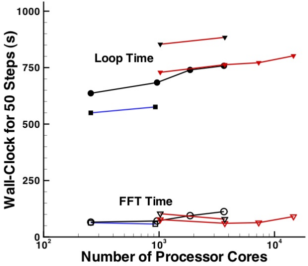

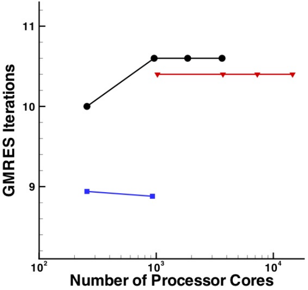

Parallel scaling challenges for peta-scale NIMROD computations

- Scaling of preconditioners
  - Scaling of direct solver (SuperLU) may become bottle-neck for current NIMROD implementation.
  - Threshold-based ILU and new hybrid versions of SuperLU have potential for better scaling.
- FFT and domain swapping
  - Collective communication may be improved by more scalable point-to-point communication?
  - Serial FFT may be improved by parallel FFT?
- Scaling on new generation of platforms
  - Multi-core system: may require mixture of MPI and OpenMP.
  - Cell and GPU systems: may benefit hybrid version with particle closures.

Summary: Challenges for NIMROD to achieve scalable petascale computations

- Scalable numerical algorithms for solving ill-conditioned large sparse matrices
  - 2D and 3D iterative solvers
  - Preconditioners
- Scalable parallel communication schemes
  - Balance between collective and point-to-point communication
  - How to take advantage of new generation of platforms
    - Sparse matrix solver for fluid part
    - Hybrid version with particle closure