What were the original goals of CEMM and near term priorities?

Model development: GV stress, 2-fluid, non-local parallel closure

M3D Code Development: C^1 code development, field aligned mesh

NIMROD Code Development: Semi-implicit algorithm for 2-fluid and GV, non-local stress closures

AMR Code Development: Flux-surface grid, improved anisotropic diffusion

Applications: 3D halo currents during disruption, Sawtooth, pellet injection, toroidal flow damping due to error field, high-n alpha-driven TAEs and alpha-particle transport, resistive wall modes, NTM, ELM

Response to suggestions by the PAC from last year:

The PAC suggests that the CEMM explore the possibility of quantitative computational tests of the explosive instability theory of nonlinear ballooning recently proposed by Wilson, [Cowley], *et. al*.

This is being explored. Cowley attended the January 06 Fluid ELM workshop in Boulder. Zhu, Hegna, and Sovinec are comparing NIMROD with a theory of line-tied g-modes. Theory extended from $\varepsilon \sim n^{-1}$ to $\varepsilon \sim n^{-1/2}$ regime. Compressibility important. No finite time singularity. Non-ideal effects become important.

We also suggest that the CEMM include some consideration of subgrid scale type models of microturbulence effects on MHD as part of its program on kinetic closures. This is being considered. Diamond and McDevitt attended Closures workshop in March 06. We may propose something in next renewal cycle. Needs more discussion.

The PAC recommends that the nonlinear sawtooth benchmarking exercise between the M3D and NIMROD be brought to completion and appropriately documented. We couldn't agree more with this. It is one of our highest priorities to get this to a state where we understand it well enough to publish it. This will happen.

Since a major component of the CEMM interactions with the SciDAC (ISICs) have come in the (AMR) area, the actual plans for integration of the AMR capability into the two major codes, M3D and NIMROD, needs to be specified. The AMR code is a block structured finite-volume AMR code and that technology is not directly transferable to NIMROD and M3D, which are Finite Element codes. However, the value of mesh refinement has been demonstrated, and we are working with TSTT (RPI) and others (Glasser) to incorporate adaptive meshing into the next generation finite element code versions.

How much computer time have you used and on what configurations?

- NERSC:
  - ~ 1.5 M hours so far this accounting period (estimate ~2 M for the year). Essentially all jobs < 1000p. Repositories mp288, mp21, mp94, m127, m275.
- ORNL:
  - Do not have exact usage statistics. We are now part of a new 3M hour allocation on the XT3. Most of the earlier effort spent on the X1E was on developing the vector elliptic solver, and did not use large amounts of time. We did not get any X1E time in the new allocation.
- Various clusters:
  - Particularly at U. Wisconsin and PPPL. Used primarily for code development, small exploratory production runs, and debugging. (~0.5 M hours total). Codes perform very well on SGI Altix.

How would you use a petascale facility?

- A “brute force” extrapolation from CDX-U to ITER would require an increase of computer power of 10^{12} as discussed in my project presentation and documented on our website w3.pppl.gov/CEMM
- Thru problem reformulation, incorporating improved (known) algorithms, and algorithmic advance, we could achieve the same physics result with hardware increases of only 10^3 (tera→peta). However, note that this is very labor intensive and implies a significant increase in code development effort.
- Other techniques, such as greatly expanding the use of parallel kinetic closures (either continuum or PIC) would also become feasible.
- Note: Extended MHD fusion codes, like most multi-physics, multiple timescale codes, will always perform best on computers with good inter-processor communication (high bandwidth, low latency). The SGI Altix is much preferred to the IBM SP3 (seaborg) for this reason. A “petascale” facility without a petascale communications network isn’t really petascale.
- Also, Extended MHD codes will always require long running times because of the multiple timescales present. More likely to use 1/10th of the machine for 10 hours than the whole machine for 1 hour.