What are the 3 most important scientific accomplishments you have made that impact predictive capability? How have they been validated against experiment?

- Self-consistent simulation of 3 sawtooth cycles
  - qualitative agreement with CDX-U (within a factor of 2)
  - more quantitative comparison in the works
- Nonlinear simulation of Edge Localized Modes (ELMs) in D-IIID
  - demonstrated poloidal localization of mode in helical bands
  - demonstrated density profile changes more than temperature
- Nonlinear simulation of fishbone instability:
  - shows strong frequency chirping consistent with many experiments
  - JET, DIII-D, NSTX, JT-60
- Pellet Injection Fueling
  - qualitative agreement with CMOD, D-IIID, JET in relative eficiencies of inside and outside launch
  - more quantitative comparison in the works

List up to 3 key contributions CSET contributors have made to your project.

- APDEC
  - collaborating in the further development of the AMR MHD code
  - implemented curvilinear (flux) coordinates, partially-implicit solves

- ITAPS
  - we are collaborating with RPI on developing adaptive unstructured grids for the M3D-C^1code
  - collaboration with R. Samulyak and J. Glimm (SB) on front-tracking / multiscale modeling of pellet ablation

- TOPS
  - NIMROD, M3D, and M3D-C^1 use the TOPS solvers. We are collaborating with them to make solver improvements.

What was the science enabled by using the leadership computing facilities ?

Each of the science accomplishments listed on slide #1 required leadership class computing facilities. Two examples:

- NIMROD nonlinear ELM calculation performed for Joule milestone
  - 344 processors of the IBM SP5 (40% of Bassi at NERSC)
  - required 9 days and 71,000 processor hours
  - 76% scaling efficiency and 262 Gflops (10% of peak)
  - higher resolution and kinetic closures could take this to the ~100 TF level

- M3D simulation of sawtooth cycle in CDX-U with 1,212,048 vertices
  - 500 wallclock hours on 384 processors on Jaguar
  - 200,000 processor hours for 500,000 timesteps
  - code exhibits good strong scaling in this regime – tradeoff between more processors and batch queue wait times
  - higher resolution (for higher S) and an energetic particle component will increase computational requirements, but also improve parallel scaling efficiency
  - ITER-class runs easily at the ~100 TF level

What new non-linear two-fluid (2F) benchmarks are planned?

- 2F magnetic reconnection in the presence of a guide field
  - extension of GEM reconnection problem by adding guide field
  - promised to DOE as FY07 milestone (see http://w3.pppl.gov/CEMM)
  - preliminary results show strong effect on suppression of fast reconnection

- 2F Nonlinear sawtooth benchmark
  - Nonlinear sawtooth calculations are being extended to better correspond to experimental conditions
  - 2F terms will be used in code-code and code-experiment

- ELM benchmark
  - NIMROD and M3D are both involved in ELM benchmark activity
  - This will proceed through nonlinear 2F stage.

Can you stabilize the high-n modes in CDX-U with two-fluid terms?

- We found that the 2F terms had a very small effect on the high-n modes for the parameters we used.

- The large (experimental values of the) thermal conductivity was much more effective at stabilizing these modes

- The study of the physics of these high-n modes is not over
  - We put it on hold, since it was not one of our milestones
  - Large thermal conductivity allowed us to continue our low-n studies, but the need for this raised many questions
  - There may be a self-consistent saturation level where the modes cause thermal transport which stabilizes the modes … Carreras/Diamond mechanism
  - To be studied further

- This may all change at higher S values

What is the path to petascale applications [for extended MHD]?

The prospects for petascale applications in the extended MHD area are very favorable. The major codes scaling efficiency improves with increased resolution and with increased complexity of the physics models (parallel closures)

- NIMROD / M3D / M3D-C^1 all utilize concurrent instances of TOPS solvers [SuperLu_dist, HYPER, GMRES, etc] as preconditioners or 2D solves
  - As these solvers scale with processor number, the codes scale even better

- PIC parallel closures will dominate the running time in many applications
  - These scale the same or better than those in turbulence codes

- Continuum kinetic parallel closures (CEL) will dominate the running time in many application
  - These demonstrate very favorable parallel scaling

- High-n MHD and energetic particle modes in tokamaks and stellarators
  - These applications require many toroidal planes, which scale exceptionally well for a finite difference code (M3D)

- APDEC has committed to bringing CHOMBO-based codes to the peta-scale (AMRMHD)