STATUS AND NEEDS FOR DOE COMPUTING:
A MISSION DRIVEN PERSPECTIVE

D. D. SCHNACK

SCIENCE APPLICATIONS INTERNATIONAL CORP.

SAN DIEGO, CA

Presented to the ASCAC Subcommittee on Facilities

25 October 2001, Washington, DC

OUTLINE

- **Mission needs for DOE computing facilities**
  - **Example: Fusion Energy Sciences**
- **Background: The development of DOE computing within the fusion program (OFES)**
- **How computing relates to the present OFES mission**
- **Computational Projects within OFES**
  - **Example: The NIMROD Project**
- **Are present DOE facilities meeting our mission needs?**
- **A vision for the direction of computing to support DOE missions**

PERSONAL EXPERIENCE WITH DOE COMPUTING

- **Continuous user of DOE computing facilities since 1973**
  - **pre-history => CTRCC => MFECC => NERSC=>???**
- **Staff member of CTRCC and MFECC at LLNL**
- **Personally developed many large scale computational models that used these facilities**
- **Project leader for the NIMROD Project**
  - **Large scale, multi-disciplinary, multi-institutional computational project within OFES**
- **I am not a computer scientist, but I must solve real programmatic problems that require advanced computing resources**

LARGE SCALE COMPUTING AND THE OFES MISSION

- ***The mission of the U.S. Fusion Energy Sciences Program is to advance plasma science, fusion science, and fusion technology -- the knowledge base needed for an economically and environmentally attractive fusion energy source.***
- **Fusion plasmas present an extreme scientific challenge**
  - **The fundamental laws of physics are known, but the collective nonlinear dynamical behavior is very complicated**
  - **Dynamics characterized by extreme separation of space and time scales, and extreme anisotropy**
  - **Science problems are as challenging as fluid turbulence**
- **Scientific uncertainties can have large impact on program costs**
- **Relevant experiments are very expensive**
- **Realistic simulation using large scale computing is cost effective**

IMPORTANCE OF REALISTIC SIMULATION

- **Advanced fusion experiment will cost ~ $2-10 B**
- **Cost proportional to volume: $\$ \sim V$**
- **Power density proportional to square of max. pressure: $P/V \sim p^2_{max}$**

  ➔ $\$ \sim 1/ p^2_{max}$ for fixed P and B (engineering constraints)
- **Physics uncertainties (eg., neo-classical tearing modes) limit max. pressure to ~ 2/3 $p_{max\text{-}theoretical}$**

Uncertainties in nonlinear physics account for ~ 1/2 the cost of advanced fusion experiment!

Predictive fluid simulation with realistic parameters has high leverage to remove this uncertainty

CENTRALIZED COMPUTING STARTED WITH FUSION

- **Fusion is a decentralized program**
- **In the 1970s it was recognized that the required simulation capability could only be obtained with a centralized computer center and high speed long distance networking**
  - **Reduced duplication of effort, eliminated institutional monopolies, facilitated collaborations**

It was revolutionary!

- **This model has been subsequently adopted by other programs and other agencies**
- **Evolved into the present DOE computing environment**

IMPORTANT FUSION PROBLEMS NOT REQUIRING SUPERCOMPUTING

- **Programmatically important, but not requiring the most advanced computing facilities**
  - ***Equilibrium*: static balance between magnetic and pressure forces in experimental geometry (solving a complicated PDE)**
  - ***Stability*: Is the equilibrium stable to small displacements? (solving a complicated eigenvalue problem)**
  - ***Transport*: diffusion of mass and energy in response to microscopic processes (solving a time dependent PDE)**

The majority of fusion computing problems do not require the most advanced supercomputers

ESSENTIAL PROBLEMS REQUIRE SUPERCOMPUTING

- **Problems whose solution is essential to the OFES mission:**
  - ***Turbulent transport*:**
    - **Computing diffusion coefficients from first principles => a fundamental understanding of the basic confinement properties of fusion plasmas**
  - ***Long time scale fluid dynamics* (extended MHD):**
    - **Nonlinear evolution of slowly evolving plasma instabilities => a fundamental understanding of consequences of operating near stability boundaries**

These problems require the largest, fastest supercomputers imaginable!

CHARACTERISTICS OF FUSION COMPUTING

- **The important fusion problems are time dependent, and therefore 4-dimensional**
- **They require many thousands of time steps**
- **Often a large, ill-conditioned linear algebra problem must be solved at each time step**
- **Parallelization can be applied to the 3 spatial dimensions**
- **Parallelization *cannot* be applied to time**
  - **Prohibited by causality**

Programmatically important fusion problems are not optimally suited to massive parallelization!

FACTORS AFFECTING FUSION COMPUTING

- **Ascendancy of massively parallel computer architecture over serial vector machines, and associated conversion costs**
- **Uncertainty in configuration of future computer architecture, and anticipated costs of additional future code conversions**
- **High cost and long time line of individual (*superhero*) code development, and general inaccessibility of resulting codes by the fusion community**
- **Increasing programmatic importance of physics problems that require extensions of the capability of existing computational tools**

NEEDS OF FUSION COMPUTING

- **The need to lower the future cost of application development, maintenance, and use**
- **The need to implement computational and programming methods that isolate architecture dependencies so that future codes will be more efficiently adapted to evolving new architectures**
- **The need to make applications codes more readily available to, and usable by, the overall fusion science community**
- **The need for advanced simulation codes to address physics problems of immediate programmatic interest**

OFES responded with Project-based Computing

PROJECT-BASED COMPUTING

- **The *old* way: Code development by individual researcher (the Superhero):**
  - **Individual or institutional monopolies**
  - **Non-portable, unreadable, inflexible, non-maintainable code**
  - **Eventually required expensive rewriting**
- **The *new* way: Code development by inter-disciplinary, multi-institutional team**
  - **Publicly available (including source), no monopolies**
  - **Extensible, modular, portable, maintainable code**

Need for centralized facilities that not only provide high quality CPU cycles, but that also facilitate project function

EXAMPLE: THE NIMROD PROJECT

The vision of the NIMROD project is to provide the magnetic fusion research community with a useful and <u>widely</u> <u>available</u> computational tool that can be applied to obtain a <u>predictive</u> <u>understanding</u> of the nonlinear dynamics of the most modern and expensive fusion experiments.

- **The NIMROD project represents a *<u>new paradigm</u>* for fusion computing because:**
  - **It seeks to solve the critical nonlinear physical problems facing the fusion program within a user-friendly, widely available, self-contained structure that allows interaction with commonly available analysis tools and experimental data bases; and,**
  - **It aims not only to provide this advanced computational capability, but also to <u>change fundamentally</u> the way that capability is developed, deployed, and maintained within the fusion program.**

THE NIMROD CODE SYSTEM

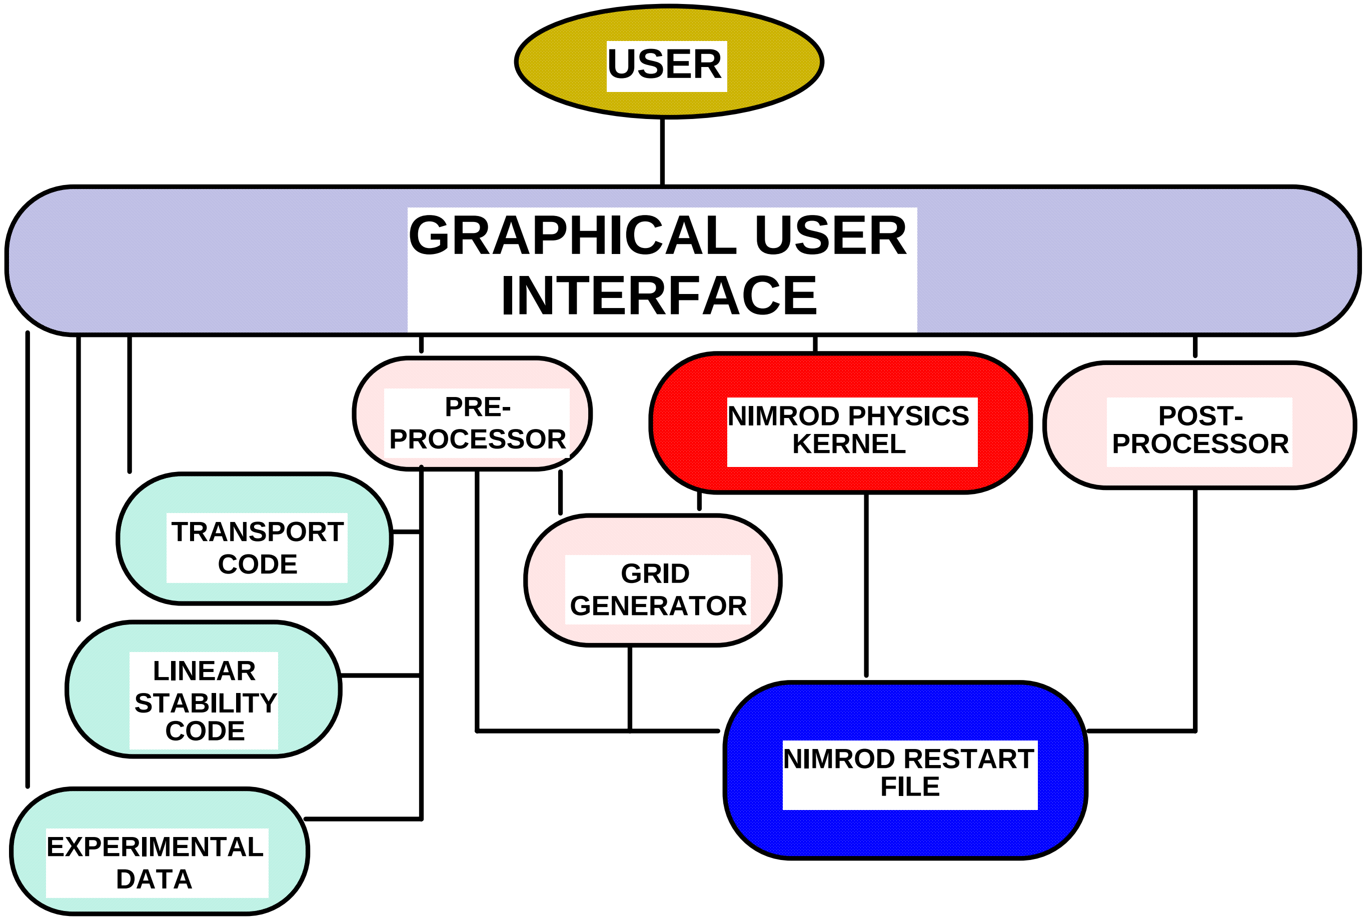

A NIMROD PHYSICS CALCULATION

- **Time-dependent, nonlinear evolution of a slowly growing, non-ideal MHD instability at large Reynolds' number ($S \sim 10^6$)**
- **Computed on T3E at NERSC:**
  - **8 dependent variables on 49,252 vertices**
  - **154 PEs**
  - **40,870 time steps**
  - **64,680 CPU hours => 7.38 CPU years**
  - **3.2 months wall clock time**
- **Increasing number of PEs has little effect because of causality**

This performance is only marginally acceptable

DIII-D SHOT 86144

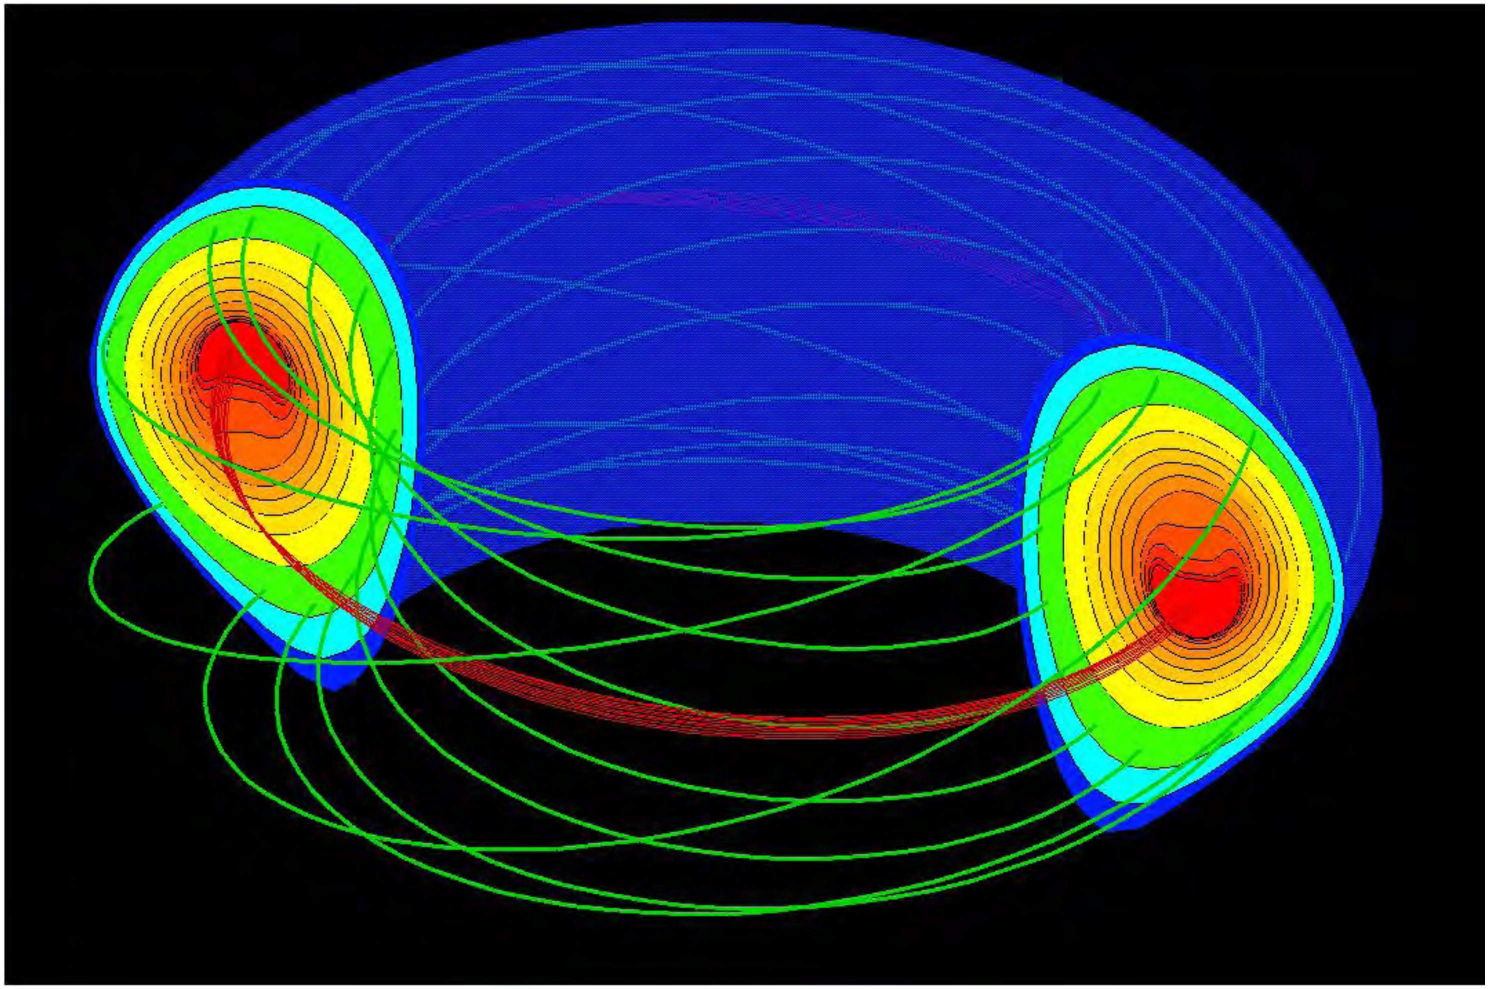

HIGH END NEEDS FOR MISSION PERFORMANCE

- **The largest, fastest systems available (can never be too big, or too fast)**
- **Optimize performance for long runs**
  - **Specialized systems for time dependent problems?**
  - **Longer residence time for jobs?**
  - **Revisit the MP paradigm?**
- **Maximize memory for large Reynolds' number calculations**
- **Offload smaller jobs to mid-range systems**
- **????**

MID-RANGE NEEDS FOR MISSION PERFORMANCE

- **Most fusion computing capacity does not require high-end computing**
- **Lack of generally available mid-range computing systems**
  - **Beowulf/Linux clusters**
  - **"Tens of dozens" of processors**
  - **Much cheaper than supercomputers, but too expensive for small groups/institutions**

Back to pre-CTRCC computing situation

- **Needed for debugging, and preparing large problems**
- **These jobs are clogging the queues on present DOE supercomputers**

NETWORKING AND OTHER SUPPORT

- **Large data "farms" for archiving data (many Gbytes/run)**
- **Fast networks for transferring data between central and local facilities**
- **Software and hardware to enable team interaction**
- **High end graphics, and the capability to use it remotely (in the real world of firewalls, etc.)**
- **Debugging and performance analysis tools**
- **Mathematical software libraries**
  - **F90. Don't insist on C++**
- **???**

VISION TO SUPPORT THE FUSION MISSION

- The largest, fastest supercomputer (of course!), but...
  - Configuration optimized for time dependent problems requiring many time steps (inherently sequential)
  - Access limited to only the largest, programmatically important problems
- Centrally available mid-range computing facility
  - Provide needed capacity for majority of fusion applications, and eliminate institutional monopolies
  - Offload jobs from the supercomputer
- Storage, networking, communication, and graphics to facilitate code and data comparisons, and enable remote collaborations

SUMMARY

- **The fusion mission is not optimally served by the MP computer paradigm**
  - **Need specialized architectures for time dependent problems**
- **Present DOE facilities are not configured in a way that well serves the fusion computing mission**
- **Need for centralized mid-range computing facility**
- **Remote collaborations are the reality**
  - **Need networks, hardware, and software that facilitates these efforts**