June 13, 2003

Professor Robert J. Goldston, Director
Princeton Plasma Physics Laboratory
Post Office Box 451
Princeton, New Jersey 08543

Dear Professor Goldston,

The Program Advisory Committee (PAC) of the Plasma Science Advanced Computing Institute (PSACI) met on June 5-6, 2003 at the Princeton Plasma Physics Laboratory. Our charge was to assess progress on the Fusion Energy Science projects within the SciDAC portfolio (five plasma science projects and the Fusion Energy Science Collaboratory), to provide recommendations for enhanced productivity, and to recommend key areas to be targeted in the next three-year phase of the Fusion SciDAC program.

The Principal Investigators (PI's) for each plasma science project and for the Collaboratory made presentations specifying the progress, the use of supercomputing resources, and the collaborations within each project and with other SciDAC activities. In each case, the PAC provided detailed feedback on both technical issues and the goals, and posed a set of questions for the PI's. These included the general questions posed by Bill McCurdy for all of the PI's: (1) What are the substantive and well-publicized scientific results obtained with the effective use of terascale computing? (2) How is the partnership with the Computer Science Enabling Technology (CSET) community delivering new capabilities for Fusion Energy Science? and (3) How does your fusion energy sciences SciDAC project play a key role in the cost-effective assessment of new ideas and operating scenarios for future experimental facilities? The PI's gave an initial response to the questions posed by the PAC and have been asked to deliver more detailed responses after discussions with their research teams. In their presentations, the PI's also detailed how the projects have responded to the previous PAC recommendations.

Assessment

Overall it is clear that significant progress has been made by each of the plasma science projects and by the Fusion Collaboratory. In the words of Marshall Rosenbluth, " A strong scientific case can be made that the SciDAC program has been a success." There have been many exciting new results enabled by high performance computing. There have also been closer collaborations among the different code groups as well as an increasingly effective interaction with the CSET community. This is likely to be one of the most important legacies of the SciDAC program. The individual scientists are to be commended.

A important general charge by the PAC to each project is that the computational and algorithmic efficiency of the codes should be very clearly identified and discussed. For example, for what fraction of the time are the computational resources being used in full parallel mode, and for what fraction of the time are computational resources being used on a small fraction of the available processors? What number of processors and amount of memory represent a good compromise between the minimum time to solution and the efficient use of system resources? What are the limitations on the performance of the codes and the concomitant restrictions upon the types of problems they can address? The PAC looks forward to hearing from each SciDAC project about these supercomputing performance issues at its next meeting.

The PAC notes that there are a number of important areas where the traditional separation of "microturbulence", "transport", and " macroscale MHD", assumed in the first phase of the SCIDAC program, breaks down. Examples include edge pedestal phenomena (ELMs, L-H transitions, separatrix effects), neoclassical tearing modes, and sawtooth crashes. The Plasma Microturbulence and Extended MHD projects may wish to consider how they could contribute to the solution of such problems, either directly in the next phase of the SciDAC program or via contributions to the Fusion Simulation project.

More detailed assessments of and recommendations for each project are now given.

1. The Extended MHD Project reported exciting new results on many topics, including high-beta disruptions in DIII-D, effects of extreme thermal anisotropy on magnetic islands, unstable toroidal Alfven eigenmodes in NSTX, interpretation of JET current-hole experiments, current bunching and ejection during magnetic reconnection, and diamagnetic stabilization of instabilities in stellarators**.** There has been an extensive use of supercomputing by this project. Interactions with the CSET community included work on adaptive mesh refinement (AMR), as well as many improvements in algorithms. A vision focussed on burning plasma issues was articulated for the next three-year program.

The PAC was pleased to see the beginnings of code comparisons on a nonlinear problem. However, the benchmarking of M3D and NIMROD for the 1-1 mode precursor to sawteeth in the CDX-U experiment has not yet led to clear conclusions. The PAC supports the new plans to bring to closure this exercise for a nonlinear sawtooth test problem, and expects to see definitive results from this comparison delivered at the next meeting. Much more emphasis on diagnosing and understanding the nonlinear dynamics in the computer simulations is strongly recommended. The M3D and NIMROD groups should move expeditiously toward regular simulation of fusion systems with a two-fluid model, since the one-fluid model is missing important physics which has a significant impact on the dynamics.

2. In the Plasma Microturbulence Project (PMP), four complementary plasma turbulence codes are now available, each having kinetic electrons as well as a common data analysis and visualization module. Three of the codes now include finite-beta effects, and one of them is being made available over the NET using the Fusion Collaboratory. Many applications are being made, including studies of ρ^* scaling and the gyro-Bohm to Bohm transition. The use of supercomputing resources is extensive.

The PAC very strongly recommends expanded comparisons among the four codes on straightforward but meaningful nonlinear problems, which are clearly tractable. We consider this to be an essential deliverable at the next meeting and request that a detailed plan for the benchmarking and enhanced collaborations be submitted for review by the PAC as soon as possible. The coordination of the four code groups is a special challenge. Closer coordination is urgently needed, including a workshop involving all four code groups on a regular (e.g., six months) basis, as well as conference calls (e.g., monthly). More detailed turbulence studies, in addition to profile and χ_i plots, are needed. Diagnostic and analysis tools, such as the common analysis software developed by Nevins, should be routinely used to reap the full benefits of the codes for physics understanding. The PAC encourages the simulation of an internal transport barrier and requests formation of a strategy for addressing in the future cross coupling between electron and ion scale turbulence. More visibly productive coupling to SciDAC CSET projects is also strongly recommended.

3. The CMRS project on magnetic reconnection is an effort connecting with the astrophysics and ASCI communities on topics of importance to fusion energy sciences. Progress was reported on the development of a magnetic reconnection code with adaptive mesh refinement (AMR) and on various applications such as the sawtooth instability. For example, a cylindrical Hall MHD code with static non-uniform grid is now being benchmarked with published results on the m=1 kink instability in a large-aspect-ratio torus. There have been close interactions with the SciDAC TOPS (Terascale Optimal PDE Simulations) on fully nonlinear implicit parallel solvers and on AMR. Since this project is still building up in many respects, it has not as yet demonstrated significant utilization of supercomputing resources.

The PAC was disappointed to find that the 3D toroidal Hall MHD code with AMR in the Flash framework, originally promised for delivery in this funding cycle, has now been deferred to the next funding cycle. A plan and schedule are requested for producing physics results using the 3D viscous MHD toroidal code and separately using the 3D Hall MHD code with AMR capability. The extent to which supercomputing resources are actually being utilized also needs to be clarified.

4. The Wave-Plasma Interaction Project is very well managed and is producing significant and well-publicized physics results utilizing terascale computing. For example, the first fully resolved 2D calculations of the conversion of fast (ion cyclotron) waves into short-wavelength modes were carried out and successfully compared with experiments. Effects of non-Maxwellian distribution functions on wave propagation and absorption are also being explored. There have been very productive collaborations with the SciDAC SSAP on parallelizing, optimizing, and restructuring the codes. Supercomputing is strongly used.

As this project has been treating linear wave propagation dynamics, the PAC strongly recommends a greater emphasis on obtaining a much more efficient code (for example, by moving from a dense matrix formulation to a sparse matrix formulation). This would enable a broader use of this capability in the fusion community and would ideally position this project for strong participation in the Fusion Simulation Project (FSP). The addition of nonlinear effects, such as self-consistent incorporation of the modified distribution functions, should also be a major priority. It is further recommended that a plan be formulated to address RF coupling with edge plasmas.

5. The Atomic Physics Project continues to deliver a wide range of results important for the diagnosis and modeling of fusion plasmas. New results have been obtained on numerous topics, such as electron-impact ionization of Li and Be, dielectronic recombination, and ion-atom collisions. Accurate atomic data is essential for a collisional radiative model. The improved calculations enabled by high performance computing are especially important for the plasma edge diagnosis at DIII-D and at JET. This project is extremely effective at widely disseminating their results and consistently generates a remarkable number of publications. This is also an activity which increases the visibility of Fusion Energy Science in other areas of physics. There has been little interaction with the CSET community so far, but plans are underway.

The PAC recommends that a very close interaction with the edge modeling community be established to ensure that the most important problems are addressed. A detailed plan to better connect with the edge modeling and diagnostics communities should be formulated. Such a plan would position this project for strong participation in edge turbulence and modeling studies in the next funding cycle.

The National Fusion Collaboratory has continued to deliver impressive technical results. The FusionGrid has been established, the TRANSP data analysis fusion code is now remotely accessible, and the GS2 microturbulence code is being implemented. Strong outreach to the user community continues. For example, large-scale demonstrations were presented at three major fusion science meetings. New capabilities in 3D visualization and animation have also been facilitated. For the ITER Project, a proposal has been developed that the U.S. take primary responsibility for software for data acquisition and management, and for remote participation. The PAC identified Steve Eckstrand as the key point-of-contact for helping to identify additional codes to be made remotely accessible. We also urge a consideration of how the Collaboratory can be used to advance graduate education in the fusion program by making available simple plasma computational tools for teaching purposes.

As its capabilities become a more established part of the fusion computing infrastructure, the Collaboratory needs to consider the issue of the most effective approach to development and support of its software. In particular, the Collaboratory should address how build versus buy decisions are made, both for immediate needs (e.g., visualization software), as well as for the long term. The Collaboratory is advised to consider whether corporate partners might at some appropriate time become interested in contributing to development of the software as its capabilities become better demonstrated and potentially applicable to a larger application domain.

The PAC notes that capacity computing resources are clearly needed to accelerate scientific productivity in the fusion energy sciences program. In view of the encouraging progress in the development of the FusionGrid by the National Fusion Collaboratory and the successful deployment of a pilot capacity computing facility at PPPL, we urge OFES to develop a compelling distributed computing strategy.

Recommendations for the Next Phase of SciDAC

The PAC was also asked to make recommendations for key areas to be targeted in the next three-year phase of the Fusion SciDAC program, which will be focussed on developing reliable computational modeling capabilities for dealing with burning plasma physics issues relevant to ITER. The PAC was unanimous in recommending three key areas which are vital building blocks for the future development of an advanced integrated modeling capability. These are 3D Extended MHD; Plasma Microturbulence; and Edge Turbulence and Modeling. The call for proposals should specify as essential parts of the workscope extensive code intercomparisons to elucidate and develop a physical understanding of these issues, a strong partnering with the CSET community, and the development of diagnostic packages to fully exploit the expanded code capabilities and gain maximum physics insight from the simulation results.

It is very important that the projects in these areas be sufficiently broad to address a spectrum of important issues. Toward this end, the following recommendations indicate illustrative topics in each of the key areas, some fraction of which (but probably not all) might be covered by proposed projects.

1. Proposals for studies of macroscale dynamics should show how the proposed research program will address relevant physics in key extended MHD scientific areas, such as a) full 1-1 mode sawtooth cycle modeling in fusion-grade plasmas, b) tearing and neoclassical tearing mode excitation and control in high-beta plasmas, c) nonlinear evolution of resistive wall modes, including toroidal flows and their control, d) effects of fast ions (e.g., alpha particles) on MHD phenomena in toroidal plasmas, e) edge MHD-type instabilities including their nonlinear evolution, and f) relevant two-fluid and kinetic effects on MHD modes in toroidal plasmas.

2. Proposals for studies of microturbulence and transport dynamics should show how the proposed research program will address the relevant physics for key microturbulence scientific questions, such as a) Bohm versus gyro-Bohm scaling and the transition between the two regimes, b) transport barrier formation and dynamics, c) the statistics of mesoscale intermittency in transport, such as avalanches, d) the dynamics of transport perturbation events, e) electromagnetic turbulence and electron heat transport due to magnetic perturbations, and f) high-beta microturbulence, as appropriate to spherical toroids.

3. Proposals in the area of edge turbulence and modeling studies should address a number of the key scientific issues, such as a) the effects of separatrix, scrape-off-layer (SOL), and x-point geometry, b) incorporation of core heat flows and edge particle sources, c) compatibility with a variety of turbulence models, such as multi-fluid electrons and ions, gyrokinetic ions and multi-fluid electrons, and kinetic

electrons and ions, d) nonlinear structures and intermittency, e) the transition to and the development of the L to H transition, f) the role of atomic and molecular physics in the structure and dynamics of the edge plasma, g) the physical effects of proximity to MHD stability boundaries and of ELM phenomena, and h) the role of plasma surface interactions. The proposals should also address how the codes will deal with the absence of a clear-cut scale separation, as in the core.

Although valuable areas, including Alpha Physics, Atomic Physics, Wave-Plasma Interactions, and other high-priority areas, could be considered important topics for funding in the next three-year phase of the Fusion SciDAC program, the three topical areas endorsed by the PAC are deemed the most vital to the success of the Fusion SciDAC program. This should be reflected in the call for proposals and also taken into account in the project selection process. As previously noted, continued work in the Atomic Physics area has the potential for strong participation in the edge physics proposals. Similarly, continued work in other areas can, as most appropriate, either be integrated into proposals in the three priority topical areas, and/or seek to participate in future FSP activities.

Finally, the PAC is grateful for the attendance and many helpful comments of SciDAC Director Alan Laub and Steve Eckstrand and Arnold Kritz from OFES. We continue to applaud William Tang and Vincent Chan for their able leadership and coordination of these many disparate national projects. We are also grateful that you were personally able to dedicate the time to attend the sessions and provide your insights and feedback.

Sincerely,

William L. Kruer
for the PSACI
Program Advisory Committee