June 14, 2006

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 at the Princeton Plasma Physics Laboratory on May 24-25, 2006. The principal charge to the PAC was to assess the progress of the three fusion energy science centers supported by the Scientific Discovery through Advanced Computing (SciDAC) program both in meeting their scientific and computational goals and deliverables as well as in implementing the recommendations made in the 2005 PSACI PAC report. The three fusion science projects are virtual centers in Extended MHD, Plasma Microturbulence, and Wave-Plasma Interactions. An additional charge was to provide initial impressions of the goals and scientific achievability of the two newly launched SciDAC Fusion Simulation Project (FSP) centers and the new "Edge Simulation Laboratory" supported by the Office of Fusion Energy Sciences (OFES) and the Office of Advanced Scientific Computing Research (OASCR). To this end, we received presentations by the Principal Investigators (PIs) of each Center. We were also given presentations by the PIs of the two Proto-FSP projects and by a representative of the newly-launched Edge Simulation Laboratory.

The presentations by the Fusion Energy Sciences SciDAC centers highlighted key technical achievements during the second year of the 3-year funding cycle. The PIs for the Fusion SciDAC centers (S. Jardin, W. Lee, and P. Bonoli) made oral presentations at the meeting and also provided two-page documents that highlighted significant research accomplishments. The PIs were requested to describe: (1) how well each project has made progress toward achieving its scientific targets with respect to clear deliverables; (2) how super-computing resources have enabled the achievement of the targeted scientific goals in a timely manner; and (3) what role collaborative interactions have played within each project and also with other SciDAC activities. The PAC's role was to provide an evaluation/assessment of substantial progress made by each project toward the scientific/computational goals and deliverables targeted by the Fusion SciDAC centers (with respect to Scientific and Technical Merit, Readiness for Terascale Computing, and Potential for Impact on Other Scientific Disciplines).

The presentations which were given to the PSACI PAC are posted on the web (http://w3.pppl.gov/theory/PSACI.html). These presentations demonstrate that the technical and computational advances have continued both to be quite substantial, and are detailed in numerous journal publications and invited papers at major meetings. The forefront advances in the modeling have also continued to be prominently featured in presentations to other scientific communities and to the funding agencies. To complement the detailed technical presentations and to focus attention on some key overarching issues, the PAC requested that each PI prepare four viewgraphs on the following topics: (1) the original project goals and deliverables, (2) the project response to the 2005 PAC recommendations, (3) the project use of terascale computing, and (4) the plans for use of petascale computing. These presentations on the second day of the PAC meeting stimulated additional valuable discussion and feedback and emphasized some very important points. Each center has made significant progress toward achievement of its goals and deliverables and, in the PAC's view, appears to be on target to achieve them at a reasonably complete level. The centers were indeed responsive to the 2005 PAC recommendations. All of the centers made productive use of the terascale computing resources with varying degrees of effectiveness. Detailed plans need to be clearly articulated as to how petascale computing will be utilized in the near future, especially for the CEMM center. New proposals for the next funding cycle are due in less than a year. To be competitive, it is essential that the next round of proposals highlight the exciting new physics discoveries relevant to ITER that are expected be achieved and how their strategic path to petascale computing will enable this to be carried out in a timely manner.

To illustrate the significant advances, let us briefly discuss some highlights in each area.

In productive partnerships with the SciDAC Computer Science and Enabling Technology (CSET) community, the Center for Magnetohydrodynamic Modeling (CEMM) has continued to make substantial advances in the development and application of improved MHD physics models. Two of the major codes (M3D and NIMROD) now include the two-fluid terms and the gyroviscous tensor, and several approaches are being pursued to develop the parallel closures in a computationally tractable form. New results were obtained on many important topics, including sawtooth behavior in the CDX-U tokamak, edge localized modes (ELMs) in the DIII-D tokamak, energetic particle modes relevant to ITER, and pellet injection for fueling tokamaks. For example, there were impressive multi-mode simulations of ELMs using the NIMROD code, as well as detailed simulations of pellet fueling with a code utilizing modern AMR techniques. While progress has been made on the nonlinear benchmarks, these exercises (e.g., the CDX-U comparison cases) need to be completed for experimentally relevant parameter regimes and documented in a timely manner. The PAC recommends that the physics applications now be focused on nonlinear problems that exercise the new 2-fluid capabilities with emphasis on discovery of new physics results enabled by the new capabilities. Progress was documented in over 20 publications, including 7 in the Physics of Plasmas and 1 in Physical Review Letters. It is also noteworthy that 9 graduate students and 4 post-doctoral researchers are being trained. Employment of terascale computing resources was exemplified by several million node-hours used on the IBM SP3 at NERSC, but most simulations were carried out in runs involving a few hundred processors or less. The PAC recommends that the CEMM scale its applications to more effectively utilize today's terascale computing and tomorrow's petascale computing. Future utilization by the CEMM of up to 10,000 processors or more on leadership-class computing platforms will depend in large measure on the successful development and implementation of more efficient implicit schemes and associated algorithms and methodologies. The exciting new science enabled by such deployment should also be clearly articulated. This is an especially important issue for the next funding cycle, when it will be necessary to show an achievable path to petascale computing.

The Center for Gyrokinetic Particle Simulation (GPS) of Turbulent Transport in Burning Plasmas has also made substantial progress. The global PIC codes now allow for shaped plasmas in general geometry with the goal of simulating ITER plasmas. The effects of velocity space nonlinearity in turbulent transport have been further studied using both adiabatic and non-adiabatic electrons in simulations of ITG turbulence. All of these simulations show similar behavior: enhanced zonal flows and a long-time steady-state flux with nearly constant amplitude. It is now generally agreed that these ITG simulations are not being influenced by particle noise. Noise effects on ETG turbulence remain under study but are beginning to seem of diminishing relevance. The PAC applauds the Center for organizing a well-attended open workshop on "long-time simulations of kinetic plasmas," and for implementing realistic collision operators into the GTC code for both neoclassical and turbulence investigations. We look forward in particular to the results from studies of the influence of collisions on ITG and TEM modes, and the turbulence and transport they induce in NSTX and DIII-D. Other physics topics to consider include electron heat transport by mechanisms other than ETG turbulence, particle transport including the physics of inward pinches, and toroidal momentum transport. The PAC recommends the development and implementation of appropriate diagnostics to deal with the complex phase-space structures in the high-resolution global simulations of turbulence. This will require more extensive interactions with the SciDAC CSET community, where this Center has productively interacted in the area of solvers. It is evident that terascale computing is being aggressively utilized in this Center. Specifically, over 4M processor hours have been used during the past year to address and help resolve particle noise issues, long-time transport evolution trends, and the kinetic dynamics associated with turbulence broadening. The codes are very portable, scalable, and efficient on both cache-based and vector-parallel MPP platforms. The GTC code, a gyrokinetic toroidal particle simulation code, which is the Fusion Energy Science representative in the NERSC benchmark suite for testing the most advanced computing platforms, provides the most impressive example of "leadership class" computing. The GTC code has now achieved 7.2 teraflops sustained performance with 4096 processors using 13.2 billion particles and with 25% single-processor efficiency on the Earth Simulator Supercomputer in Japan. The Center is well-positioned to take advantage of petascale computing but will need to clearly articulate its strategic vision on how this will accelerate progress on key turbulent transport physics discoveries. About 14 journal articles have been published or submitted for publication, and about 5 invited presentations have been given or are scheduled to be given by this SciDAC Center.

The Center for Wave-Plasma Interactions has made commendable progress on the development of an integrated simulation capability for the ion cyclotron, lower hybrid, and electron cyclotron range of frequencies including self-consistent wave fields and plasma velocity distributions. Recent scientific accomplishments include the first ever self-consistent ICRH simulations, a new understanding of the role of finite drift orbit effects in ICRH experiments on the Alcator C-Mod and DIII-D tokamaks, and identification of a possible ICW mode conversion scenario for ITER. The PAC commends these continuing productive interactions with the experimental community. This Center was especially responsive to the 2005 PAC recommendations and has, for example, developed a diagnostic to assess the validity of quasi-linear theory. Moving toward the future beyond the current compute-intensive linear dense-matrix inversion of the wave equation, attention is now being given to the important challenge of developing more advanced wave solvers using a multi-scale adaptive spectral representation of the fields. The PAC commends this thrust and looks forward to hearing about progress on this key aspect of the project. With regard to the use of terascale computing resources, it was reported that over 10^5 processor hours per month were used by this Center during the past year – an amount that is considerably less than the several million hours used annually by both the CEMM and GPSC Centers. This SciDAC Center intends to target the use of petascale computing to enable 3D ICRF heating and mode conversion studies and to do full wave LH field solutions for ITER applications. The Center's important progress has been documented in approximately 5 journal publications.

A number of specific findings and recommendations pertaining to the OFES PSACI Centers are summarized below:

1. The PAC is pleased that OFES has taken steps to make effective capacity computing more available. This is enhancing both the development of codes and their timely application to important physics problems.

2. The PAC continues to emphasize from a verification and validation perspective that more in-depth physics analysis of the simulation data and more extensive quantitative comparisons with analytic theories and with experimental observations would significantly increase the impact of the SciDAC centers on the fusion program. To this end, the PAC recommends that all of the SciDAC projects give more attention to the development and implementation of modern diagnostics and analysis tools in addition to code development. It is also clear that more collaborations with fusion theorists and experimentalists which are focused on verification and validation of the simulation results would be beneficial to the SciDAC projects. To this end, OFES should consider making resources available to facilitate such collaborations. In other major computational science programs (such as the prominent DOE NNSA-sponsored ASCI Program), substantial verification and validation efforts demand significant commitments from key experimental projects to provide the associated run time as well as the enabling infrastructure costs.

3. The PAC understands that due in good measure to the impressive achievements of the National Fusion Collaboratory project led by D. Schissel, which was previously supported entirely by OASCR until the summer of 2006, the next phase of the SciDAC Program (SciDAC-2) will provide $2M per year in funding

coming jointly from OFES and OASCR over the next five years to support this key area. Proposals are currently under review with an outcome that will hopefully enable OFES to continue to play a major leadership role in building collaborative infrastructure for future remote operation of experiments and in fostering collaborative development and use of simulation codes as well as data analysis codes. From a verification and validation perspective, this vital area of applied research is expected to help establish common data structures and protocols and a collaborative culture for integrating results from the OFES SciDAC projects (including the new proto-FSP's) into the broader fusion community, and in particular the large experimental communities both internationally and nationally. The PAC looks forward to learning next year about the goals and strategic targets for such a newly-launched OFES SciDAC science application center with strong collaborative linkages to experiments.

The PAC also looks forward to hearing the progress next year of the two recently launched SciDAC Proto-FSP Centers -- the Center for Plasma Edge Simulation (CPES led by C. S. Chang) and the Center for Simulation of Wave Interactions with MHD (SWIM led by D. Batchelor) -- and of the new Edge Simulation Laboratory (led by R. Cohen), which is supported by OFES and OASCR. The roadmaps presented to the PAC on the Proto-FSP Centers appeared to be reasonable approaches. We commend OFES for launching two edge physics projects – an extremely challenging but important topical area about which the PAC has been concerned for a number of years. In particular, we would like to hear about progress in turbulence physics results from CPES as well as how the continuum approach (presented to the PAC by J. Candy) is impacting substantial research progress in the new Edge Simulation Laboratory. We believe that it is a practical approach that the SWIM Project (as described by L. Berry) has initially implemented by using simple linkages among codes to rapidly get started on the project. However, we do encourage strong interaction with the SciDAC CSET community (within this Center as well as in other SciDAC Enabling Technology Centers and Science Application Centers) to incorporate more elegant integrative framework solutions and common data standards on the longer time scale.

In summary, the ongoing achievements of the OFES SciDAC projects make clear that advanced computations in combination with theory and experiment provide a powerful tool for scientific understanding and innovation in OFES research. Plasma science is indeed effectively utilizing the exciting advances in information technology and scientific computing, and tangible progress is being made toward more reliable predictions of complex properties of high temperature plasmas. Very importantly, the OFES SciDAC projects continue to bring together physicists, applied mathematicians, and computer scientists in close and productive working relationships, which provides an excellent model for future research.

Finally, the PAC commends the OFES and in particular Dr. John Mandrekas for the very productive benchmarking studies which he reported. We are also grateful that Dr. Anil Deane of the Office of Advanced Scientific Research (OASCR) was able to meet with the Committee and provide important insights into OASCR's priorities and expectations from the OFES SciDAC portfolio for which it has shown great interest and associated funding support. We also thank you for your active participation in this year's meeting and for your continued strong support. Finally, the PAC again warmly applauds Dr. William Tang and Dr. Vincent Chan for their ongoing and very effective leadership of the PSACI. Thanks to their vision and strong advocacy, the fusion science community is playing a highly visible and productive role in the national SciDAC program.

Respectfully for the PSACI-PAC,

William L. Kruer
Chair, PSACI-PAC