July 12, 2007

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 June 7-8, 2007. 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 and the two SciDAC FSP (Fusion Simulation Project) proto-type centers in meeting their scientific and computational goals and deliverables. The three fusion science projects are virtual centers in Extended MHD, Plasma Microturbulence, and Wave-Plasma Interactions, and the two FSP proto-type projects are virtual centers in Plasma Edge Simulation (CPES) and in Simulation of Wave Interactions with MHD (SWIM). Additional charges were to assess progress and plans for a new SciDAC proto-FSP center for Framework Application for Core-Edge Transport Simulations (FACETS) supported jointly by OFES and OASCR and for the Edge Simulation Project (ESL) supported by OFES and OASCR. A final charge was to provide impressions of the new FSP initiative. To these ends, we received presentations by the Principal Investigators (PIs) of each center, by a representative of the Edge Simulation Laboratory, and by Arnold Kritz on the recent FSP workshop.

The presentations by the Fusion Energy Sciences SciDAC centers highlighted key technical achievements during the third year of the 3-year funding cycle, and the presentations by the proto-FSP centers highlighted progress during the second year of a planned three-year funding cycle. The PIs for the Fusion SciDAC centers (S. Jardin, W. Lee, and P. Bonoli) and for the proto-FSP centers (C.S. Chang and D. Batchelor) 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 leadership-class computing resources have enabled the achievement of the targeted scientific goals as well as demonstrated the scalability of the science with the number of processors on leadership class computing platforms; 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 very substantial and to be 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. There are a growing number of impressive comparisons with and elucidations of experimental results. 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 three most important scientific accomplishments and how they have been validated against experiments, (2) the science enabled by the use of leadership computing facilities, (3) key contributions of the CSET collaborators, and (4) the deliverables not yet attained and why. The presentations on the second day of the PAC meeting stimulated additional valuable discussion and feedback and emphasized some very important points. In collaboration with the SciDAC Computer Science and Enabling Technology (CSET) community, each of the three fusion energy science centers has demonstrated impressive advances in the ability to model key aspects of fusion experiments. In the PAC's view, the targeted goals and deliverables will be successfully accomplished in nearly all cases. All of the centers made productive use of the terascale computing resources with varying degrees of effectiveness. The PAC continues to urge that more detailed plans be clearly articulated as to how petascale computing will be utilized in the near future, especially for the CEMM center. The two proto-FSP centers are making rapid progress in interfacing different state-of-the-art codes with one another towards the goal of a more integrated modeling of future experiments. This progress is being accelerated by close interactions with the fusion energy science centers and with the CSET community.

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

The Center for Extended Magnetohydrodynamic Modeling (CEMM) has been quite productive. In response to a long-standing request from the PSACI PAC, this project has now reported very notable success in carrying out a detailed M3D/NIMROD nonlinear benchmark calculation of three complete sawtooth cycles in a small tokamak (CDX-U). CEMM also completed high resolution nonlinear simulations of Edge Localized Modes (ELM's) in DIII-D, which successfully met an OFES Annual Program Performance Target. Results from these simulations showed both poloidal localization of the modes in helical bands and the important feature that the density profile changes more than the temperature. The inclusion of two-fluid effects in these nonlinear simulations reflects important progress in treating the complications associated with dispersive waves. Other impressive accomplishments include nonlinear simulations of the fishbone instability

Deleted: Other

showing strong frequency chirping consistent with experiments at JET, DIII-D, NSTX and JT-60 as well as calculations of pellet injection fueling which were in qualitative agreement with the relative efficiencies of inside and outside launch observed in experiments on C-MOD, DIII-D and JET. This close interface with experiments is very commendable and important for code validation. It was reported that the codes now have the full gyroviscous tensor and the two fluid terms incorporated, and progress continues on including kinetic parallel closures. Code verification activities have included comparisons with analytic results on a variety of problems (the gravitational mode, hot particle modes, and linear ELM's). Both NIMROD and M3D have improved their parallel scaling efficiencies so that they now effectively use 4000-5000 processors. However, it should be noted that since this corresponds to scaling with respect only to the toroidal direction (i.e., only the toroidal mode number n without an associated increase in poloidal harmonics), the challenge to demonstrate new terascale-enabled physics results remains. There is an ongoing productive interface with CSET contributors on algorithm improvements. For example, the codes use TOPS solvers, which continue to be improved. The many results of CEMM have been reported in about 50 peer-reviewed and IAEA papers. CEMM is also making an important contribution to graduate education. Ten graduate students currently participate in the activities of the center. The PAC strongly encourages extension of the nonlinear two-fluid benchmarks to a broader range of problems. We also urge that more attention be given to articulating the path from truly terascale to petascale applications.

The Center for Wave-Plasma Interactions (WPI) also made quite impressive progress. Especially notable was the self-consistent simulation of the nonlinear evolution of a nonthermal ion distribution function created by ICRH heating. The development of synthetic diagnostics enabled detailed comparisons with measurements of the energy distribution of the minority ion tail in C-MOD experiments. Extensions are underway to include multiple ion species. ICW/IBW mode conversion in present day tokamaks and in ITER can now be realistically simulated. Indeed, the first simulations of ICW mode conversion in ITER were carried out using 4096 processors on the CRAY XT-3. Full-wave EM field simulations of LH waves, including nonthermal electron distribution functions, are being validated by comparisons with the hard x-ray emission and driven currents measured in recent experiments. A capability has been developed for including the full ICRF antenna spectrum and realistic antenna geometry, and a self-consistent treatment of the RF antenna-edge plasma is under development. Leadership class computing resources were effectively used. For example, a calculation of fast wave current drive in ITER utilized 2048 processor cores for 8 hours on the Cray XT3/XT4. Close collaboration with CSET and SAPP partners has resulted in code optimizations, improved algorithms for matrix inversion, and 3D visualizations of nonthermal particle distributions and time-dependent wave fields. These many results were reported in about 8 publications this year. The PAC applauds the development of synthetic diagnostics and the detailed comparisons with experiments as well as the progress in including the self-consistent plasma response and antenna and plasma sheath effects. For improvements, we continue to strongly encourage the development of multi-scale adaptive spectral algorithms and also recommend consideration of how to address various 3D effects, such as wave propagation with large 3D magnetic island structures or in 3D confinement systems such as stellarators.

The Center for Gyrokinetic Particle Simulation (GPSC) of Turbulent Transport in Burning Plasmas continued to make significant progress. The main activities have been centered on two gyrokinetic PIC codes (GTC-S for shaped plasmas and GTC for circular plasmas) as well as on a wedge code (GEM) for shaped plasmas. Interfaces have been implemented between GTC-S and GEM and the TRANSP code (for experimental plasma profiles) and the JSOLVER code (for MHD equilibria). These interfaces will facilitate comparisons with experiments in order to validate the turbulence codes. Two of the codes (GTC-S and GTC-Neo) have been used to successfully simulate significant features of NSTX discharges. For example, the steady-state neoclassical ion thermal diffusivities calculated by GTC-Neo are much closer to the experimental data than are the predictions of NCLASS, a commonly used code in the fusion community. The various simulations by this center have illustrated the importance of velocity space nonlinearity for the simulation of large plasmas, various nonlocal effects resulting from finite orbit effects, turbulence spreading in the radial direction, and energy cascade to lower (m,n) modes. Simulations and theory clarified the role of discrete particle noise in PIC simulations of the turbulence, and a resetting scheme resolved the issue of growing weights in the delta f scheme for specific cases. The Center has continued to provide high visibility for fusion energy science applications with its aggressive and effective use of leadership class computing. For example, a GTC simulation of ETG turbulence used 40 billion particles and 6400 processors on the CRAY XT3 at ORNL. The GTC-S code has been chosen as the only OFES code to be part of the Joule Applications, a software effectiveness exercise for OMB, and as one of the six early applications on the 250TF Cray at ORNL. The many results were reported in over 30 peer-reviewed journal articles this past year and were also featured in the SciDAC Review, Vol.1 (2006). A 3D visualization of a twisted mesh nonlinear turbulence structure used in the GTC-S shaped-plasma simulations was featured on the cover of the 2006 Battelle Annual Report. The PAC applauds the implementation of electromagnetic effects in the GEM code but is concerned with the absence of similar progress with respect to GTC-S. This effort should have the highest priority. We look forward to the results from studies of the influence of collisions on ITG and TEM modes, and actual predictions for the characteristics and level of turbulence and transport they are expected to 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 continues to recommend the development and implementation of appropriate synthetic 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, with which this Center has already productively interacted in a number of areas.

The proto-FSP Center for Simulation of Wave Interactions with MHD (SWIM) brings together two fields of fusion computation: high power wave-plasma interactions and extended MHD. The goal is to understand how RF waves can be used to mitigate and control macroscopic instabilities, such as the sawtooth and the neoclassical tearing mode

(NTM). The main accomplishment to date is in the computer science area: development of a computational framework and infrastructure for coupling codes from the related fusion energy science PSACI centers (WPI and CEMM). A first version of the so-called Integrated Plasma Simulator has been designed and implemented, and some initial coupled TSC/TRANSP runs have been carried out. Close collaboration with CSET contributors is ongoing on topics ranging from solvers to data transfer. The principal goals for next year are to have the Integrated Plasma Simulator running on two different platforms with state of the art capabilities which will allow regular users to begin studies of RF effects on sawtooth behavior. The PAC noted that there were a number of computational frameworks for code integration under development and strongly encouraged good communication with the other efforts. At this point in time, it seems both reasonable and valuable to explore various paths.

The proto-FSP Center for Plasma Edge Simulation (CPES) aims to develop a predictive capability for edge plasmas, including the pedestal equililbrium and associated transport dynamics as well as the ELM crash and the resulting heat load on the wall. Two new particle codes are being developed for edge kinetics: XGC0 and XGC1. XGC0 is a time-dependent equilibrium code which is already in production mode. XGC1 is a 3D turbulence code, whose electrostatic turbulence capability is being verified. These codes are being coupled to linear and nonlinear ELM codes for simulation of pedestal growth, ELM crash, scrape-off transport, and divertor heat load. The coupled simulations are carried out using the SDM Kepler framework. New physics results to date include a gyrokinetic 2D equilibrium solution across the toroidal magnetic separatrix in the edge plasma and a demonstration of a drop in pedestal density but not temperature by 3D resonance magnetic perturbations. This Center will clearly be an aggressive user of high performance computing, having already used up to 16,000 Jaguar cores and 1.1 million processor-hours. The Center includes five post-docs and nine graduate students and actively collaborates with other fusion energy science PSACI centers (CEMM and GPSC). There is also close collaboration with the CSET community on solvers, algorithms, and the Kepler framework. The PAC suggested that existing two fluid EM codes be considered for interim use until the full kinetic codes come online. Further given the broad recognition of the importance of electromagnetic effects in edge turbulence, the PAC encourages the team to begin upgrading the electrostatic kinetic model as soon as is practical.

The Edge Simulation Laboratory (ESL) is a project to develop a gyrokinetic code for edge plasmas based on continuum techniques. This is a base program activity of OFES, coupled with an allied algorithm research activity funded by the OASCR base Math Program. ESL's continuum approach is complementary to the PIC-based methods of the CPES project. ESL has three components: TEMPEST, a full-f, full-geometry code; EGK, a simple geometry, rapid prototype code; and a math component which develops algorithms for a next generation code. This past year TEMPEST has been further developed; for example, it has been extended from 2D to 3D. Code applications focused on verification of the 2D code in the area of neoclassical transport and geodesic acoustic modes. EGK has also undergone verification tests with comparisons to theory and other codes, such as GS2. The code recently obtained initial solutions from the 2D gyrokinetic

Poisson equation to obtain the neoclassical electric field including poloidal variation with kinetic electrons and ions. The PAC strongly recommends that a systematic verification of Neoclassical theory in the banana regime be carried out as an initial demonstration of the physics fidelity of the codes. The applied math activity has developed and documented suitable discretizations for the gyrokinetic Poisson equation and has carried out test problems. Significant progress of the ESL has been documented in about six recent publications.

Finally, the FACETS project was funded in January of 2007 as part of the SciDAC portfolio in OFES to provide a framework application for core-edge transport simulations. The goal is to provide a coupled core-edge-wall computational capability to the fusion community with maximal re-use of existing software. This project appears to be well organized and on its way to developing the framework and assembling or developing the various codes and models for dealing with the core sources, the edge, the wall, and the turbulence. There also appears to be a close coupling with the applied math and computer science partners. The PAC underscores the importance of frequent communication with the other projects developing frameworks. We note that the concept of surfacial coupling is intriguing, but it seems unclear how matching on various surfaces can properly account for complex events, such as ELM's. We look forward to seeing the progress in this area.

The final charge was to comment on the plans for the new FSP initiative. Arnold Kritz gave a clear and comprehensive presentation of the recent FSP workshop and articulated a compelling case for proceeding with this initiative. As discussed in an earlier report, the PSACI PAC is strongly supportive of the Fusion Simulation Project. A number of associated points which need to be stressed are as follows:

1. It is important to articulate the 5-10-15 year goals to be both challenging and achievable. In 5 years an integrated model with significantly improved predictive capabilities should be obtained and used to help design and interpret fusion experiments. The subsequent 10- and 15-year goals are to develop an increasingly well-validated model with ever improving physics fidelity, resulting in an unprecedented tool for designing optimal fusion devices.
2. It is very important to connect with and build on the expertise being developed in the SciDAC fusion energy science centers and the SciDAC proto-FSP centers. The "physics modules" for the integrated model need ongoing development in order to finally achieve the goal of an integrated model with high physics fidelity.
3. It is very important that the simulations be verified and validated with experiments and theory. Hence, a strong base program on fusion energy science must be engaged to help ensure a positive outcome for this aggressive initiative.
4. It is important that there be strong community buy-in on the FSP project. The advocates for FSP must clearly make a persuasive case to the community that it can provide a major impetus to pushing all fusion science forward.

A final recommendation is summarized below:

> 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 very beneficial to the SciDAC projects. To this end, OFES should consider making dedicated 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.

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 is grateful to Dr. John Mandrekas of OFES and Dr. Mark Sears of OASCR for their participation in this year's meeting and for their valuable input. We also thank you for your interest and for your continued strong support. Lastly, 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