

Princeton Plasma Physics Laboratory

PPPL-

PPPL-



Prepared for the U.S. Department of Energy under Contract DE-AC02-09CH11466.

Princeton Plasma Physics Laboratory

Report Disclaimers

Full Legal Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Trademark Disclaimer

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors.

PPPL Report Availability

Princeton Plasma Physics Laboratory:

<http://www.pppl.gov/techreports.cfm>

Office of Scientific and Technical Information (OSTI):

<http://www.osti.gov/bridge>

Related Links:

[U.S. Department of Energy](#)

[Office of Scientific and Technical Information](#)

[Fusion Links](#)

Scientific and Computational Challenges of the Fusion Simulation Program (FSP)

William M. Tang

Princeton University, Plasma Physics Laboratory

P.O. Box 451, Princeton, New Jersey 08543. USA

This paper highlights the scientific and computational challenges facing the Fusion Simulation Program (FSP) – a major national initiative in the United States with the primary objective being to enable scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration. This requires developing a predictive integrated simulation capability for magnetically-confined fusion plasmas that are properly validated against experiments in regimes relevant for producing practical fusion energy. It is expected to provide a suite of advanced modeling tools for reliably predicting fusion device behavior with comprehensive and targeted science-based simulations of nonlinearly-coupled phenomena in the core plasma, edge plasma, and wall region on time and space scales required for fusion energy production. As such, it will strive to embody the most current theoretical and experimental understanding of magnetic fusion plasmas and to provide a living framework for the simulation of such plasmas as the associated physics understanding continues to advance over the next several decades. Substantive progress on answering the outstanding scientific questions in the field will drive the FSP toward its ultimate goal of developing the ability to predict the behavior of plasma discharges in toroidal magnetic fusion devices with high physics fidelity on all relevant time and space scales. From a computational perspective, this will demand computing resources in the petascale range and beyond together with the associated multi-core algorithmic formulation needed to address burning plasma issues relevant to ITER – a multibillion dollar collaborative experiment involving seven international partners representing over half the world’s population. Even more powerful exascale platforms will be needed to meet the future challenges of designing a demonstration fusion reactor (DEMO). Analogous to other major applied physics modeling projects (e.g., Climate Modeling), the FSP will need to develop software in close collaboration with computer scientists and applied mathematicians and validated against experimental data from tokamaks around the world. Specific examples of expected advances needed to enable such a comprehensive integrated modeling capability and possible “co-design” approaches will be discussed.

Keywords: magnetic fusion energy, physics integration, high-performance computing, predictive simulations, experimental validation.

1. Introduction

A computational initiative called the Fusion Simulation Program (FSP), which will be led by the U.S. Department of Energy (DoE) Office of Fusion Energy Sciences (OFES) in partnership with the Office of Advanced Scientific Computing Research (OASCR) is currently being developed with the primary objective being to produce an experimentally-validated predictive simulation capability that is important to ITER and relevant to major current and planned toroidal fusion devices. This is expected to be a “transformational” capability that will demand the development over the next decade of advanced software designed to use leadership class computers (at the petascale and beyond) for carrying out unprecedented multi-scale physics simulations to provide information vital to delivering a

realistic integrated fusion simulation modeling tool. Modules with much improved physics fidelity will enable integrated modeling of fusion plasmas in which the simultaneous interactions of multiple physical processes are treated in a self-consistent manner. The associated comprehensive modeling capability will be developed in close collaboration with experimental researchers and validated against experimental data from the major U.S. facilities and also internationally. Since each long-pulse shot in ITER is expected to cost over \$1M, this new capability promises to be a most valuable tool for discharge scenario modeling and for the design of control techniques under burning plasma conditions.

2. Background

Fusion energy research has historically been a

leading high performance computing (HPC) application domain. Significant progress in both particle and fluid simulations of fine-scale turbulence and large-scale dynamics in magnetically-confined plasmas have been enabled by the combination of access to powerful supercomputing resources at the petascale level together with innovative advances in analytic and computational methods for developing reduced descriptions of physics phenomena spanning a huge range in time and space scales. In particular, the fusion community has made excellent progress in developing advanced codes for which computer run-time and problem size scale well with the number of processors on massively parallel machines. A good example is the effective usage of the full power of multi-teraflop to petaflop systems to produce three-dimensional, general geometry, nonlinear particle simulations which have accelerated progress in understanding the nature of plasma turbulence in fusion-grade high temperature plasmas. These calculations, which typically utilize multi-billions of particles for thousands of time-steps, would not have been possible without access to such powerful modern supercomputers together with advanced diagnostic and visualization capabilities to help interpret the results. This is significant because one of the central plasma physics problems on the road to creating a working fusion power plant is understanding, predicting, and controlling instabilities caused by unavoidable plasma inhomogeneities. One consequence is the occurrence of turbulent fluctuations (“microturbulence”) which can significantly increase the transport rate of heat, particles, and momentum across the confining magnetic field – an overall effect that severely limits the energy confinement time for a given machine size and therefore the performance and economy of a tokamak device [1, 2]. Accelerated progress on this critical issue is especially important for ITER because the size and cost of a fusion reactor is determined by the balance between such loss processes, and the self-heating rates of the actual fusion reactions [3]. Petascale computational modeling and simulation is irreplaceable in dealing with such challenges because of the huge range of temporal and spatial scales that must be taken into account [4]. Existing particle-in-cell (PIC) techniques [5, 6] have clearly demonstrated excellent scaling on current terascale leadership class supercomputers. For example, the Gyrokinetic Toroidal Code (GTC) [7, 8] a mature PIC code has demonstrated excellent scaling on virtually all of the current leadership class facilities worldwide utilizing MPI and OpenMP [9, 10, 11, 12, 13].

In general, reliable predictive calculations of turbulent transport can only be achieved through experimentally validated simulations, as it fills the space

between empirical observations, which are expensive and difficult to acquire, and theory, which cannot handle the nonlinearity of the problem under realistic experimental conditions. To do this, the majority of researchers world-wide use a gyrokinetic (ab initio) approach. For example, PIC codes solve the nonlinear equations underlying gyrokinetic theory with excellent scaling to more than 100,000 cores having already been demonstrated. However, in order to move in a timely manner to producing HPC simulations with the highest possible physics fidelity, it is expected that computing at the exascale will be necessary. The associated software/algorithmic advances needed must be compatible with the evolving hardware and therefore developed in a true “co-design” sense. With the unprecedented resolution in a multi-dimensional phase-space enabled by access to HPC platforms at the petascale and beyond to the exascale, these advanced kinetic simulation capabilities are expected to have direct relevance to existing experimental devices as well as to ITER. The impressive advances achieved by such codes give great promise of delivering scientific discoveries appropriate for “path to exascale” grand challenges.

From the HPC perspective, the success of the PIC codes provide important “lessons learned” of how modern algorithmic approaches can expedite progress toward needed new insights [9]. Motivated by the strong shift to multi-core environments extending well beyond the quad-core level, PIC studies have focused on improving current programming frameworks — such as systematically testing the proposed two-level hybrid programming model. In making full use of the multiple cores on a node, it is necessary to avoid being constrained to an MPI process on each core. Since some arrays get replicated on all these processes, a memory limit is encountered for the larger problem sizes of interest. The “Hybrid Programming” method provides a much more viable way to mitigate this problem. Results supporting the efficacy of such an approach have been successfully obtained on the quad-core Leadership Class Facilities — Intrepid (IBM Blue-Gene-P) at Argonne National Laboratory (ALCF) and Jaguar CRAY XT-4 at Oak Ridge National Laboratory (OLCF). As a specific example, the full deployment of the more efficient hybrid OpenMP/MPI algorithm together with a new radial domain decomposition feature in the GTC-P code operating on the Intrepid LCF at ANL. This capability greatly facilitates examining the key question of how plasma microturbulence properties might be affected as the plasma size increases from that of existing experiments to the future very large plasmas characteristic of ITER. In particular, it is now possible to efficiently examine the global turbulence properties in

devices ranging from current scale experiments, which exhibit an unfavorable “Bohm-like” scaling with plasma size to the ITER scale plasma which is expected to exhibit a more favorable “gyro-Bohm” scaling of confinement. Multi-scale global kinetic simulations of plasma microturbulence, which take into account both the micro and meso scales of interest, can systematically investigate the turbulent transport properties of a magnetically-confined thermonuclear plasma spanning the size of existing machines today to the ITER device which will be more than 3 times larger in minor radius. Gaining better understanding of the expected improved confinement for these much larger ITER-scale plasmas is critically important because the microturbulence-driven nonlinear diffusion of particles and energy works against the plasma self-heating produced by the fusion reactions. The “scientific discovery” aspect of such studies is that while the simulation results can be validated against present-day tokamaks, there are no existing devices today that have a minor radius that is even one-third the size of ITER. Accordingly, the role of high physics fidelity predictive simulations takes on an even more important role — especially since the expected improvement in confinement for ITER-sized devices cannot be experimentally validated until after it is constructed and operational. Such advanced kinetic simulations also help illustrate how computational capabilities at the extreme scale can expedite the delivery of key scientific insights appropriate for grand challenges at the petascale and beyond. The “weak” scaling of GTC-P (>130,000 processors), which is enabled by the parallel algorithm implemented with MPI and OpenMP and developed specifically for the IBM-Blue Gene-P system at the ALCF, is illustrated below:

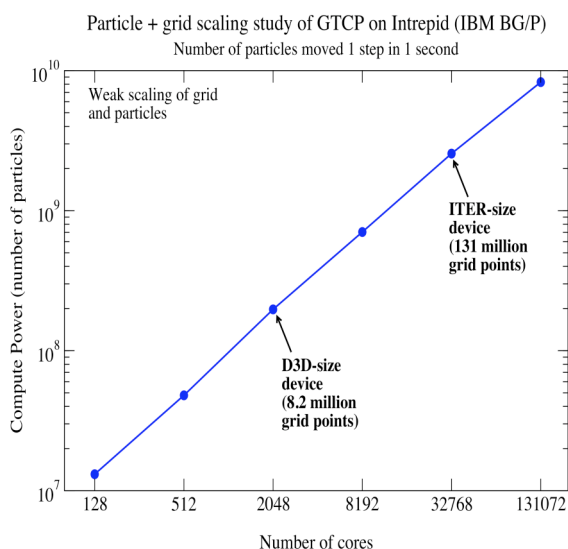


Fig 1: Speedups of the GTC-P code on the quad-core IBM Blue-Gene-P (results courtesy of S. Ethier)

3. FSP Goal

The goal of the FSP is the development of a suite of modern software tools that are experimentally validated to enable reliable predictive simulations for magnetically confined fusion plasmas in regimes and geometries relevant for practical fusion energy. In accomplishing this objective, the FSP will leverage the emergence of high performance computing (HPC) resources at the sustained petascale level and beyond together with the scientific programs from the U.S. Department of Energy (DoE) Fusion Energy Sciences (FES), Advanced Scientific Computing Research (ASCR), and especially those under the auspices of the Scientific Discovery through Advanced Computing (SciDAC) Program, which has demonstrated over the course of the past decade the importance of effective interdisciplinary collaborations between computer science, applied mathematics, and many domain applications areas, including Fusion Energy Sciences. The FSP is expected to: (i) be an important asset for supporting the scientific mission of ITER; (ii) enable more efficient harvesting of physics insights from major current toroidal fusion devices; and (iii) provide new integrated modeling tools for exploring future magnetic confinement systems, including an eventual demonstration reactor (DEMO). The associated major challenge is to develop a suite of advanced software codes designed to effectively utilize leadership class computers for carrying out multi-scale multi-physics simulations that deliver a realistic integrated fusion simulation capability with unprecedented fidelity in physics models, solution accuracy, and geometric representation. This will demand close collaborations with experimental researchers to validate the new software tools against data from magnetic fusion energy (MFE) facilities — both nationally and internationally. Experimental validation involving the deployment of modern diagnostics is essential to this vision for producing what will be an embodiment of the theoretical and experimental understanding of magnetically-confined thermonuclear plasmas and will provide a “living framework” for the simulation of such plasmas as the associated physics understanding continues to advance/evolve over the next several decades. Substantive progress on answering the outstanding scientific questions in the field will drive the FSP toward its ultimate goal of developing a reliable integrated ability to predict the behavior of plasma discharges in toroidal magnetic fusion devices on all relevant time and space scales.

4. FSP Benefits

A successful FSP will enhance the return on investments in fusion experiments in general and greatly

facilitate the harvesting of scientific insights from ITER. This in turn will hopefully enable discovery of new modes of operation with possible extensions of performance enhancements and improvements needed for DEMO. High fidelity, whole-device modeling capabilities in Fusion Energy Sciences will demand computing resources in the petascale range and beyond to address ITER burning plasma issues. Sustained petascale and even exascale (10^{18} floating point operations per second) platforms (likely to appear circa 2020) will be needed to meet the future challenges of designing DEMO.

Overall, the FSP will embody our state of knowledge in a suite of advanced codes under a unified framework and made widely available to the FES community. With regard to some unique benefits for the FES area, potential FSP contributions include:

- Addressing multi-physics and multi-scale problems that are now treated in isolation, leading to scientific discovery of new phenomena that emerge only with integration;

- Carrying out a rigorous and systematic validation program in collaboration with experiments to put models on the firmest and most realistic possible foundation;

- Identifying a set of Science Drivers (SD) as forefront scientific problems in FES and developing a “living roadmap” to address them as Integrated Science Applications (ISA’s) that will be aided by computing at the extreme scale in collaboration with ASCR over the next decade;

- Developing predictive models which improve our capabilities for reliable scenario modeling, especially for ITER, and for the design of future machines such as DEMO;

- Incorporating powerful HPC capabilities to help accelerate scientific understanding and modern software engineering approaches to ensure the reliability, robustness, and ease-of-use of the new tools that are developed; and

- Leveraging ongoing activities (theory, experiment, modeling in FES and applied math, computer science in ASCR) to develop unprecedented simulation capabilities.

With regard to unique benefits for Advanced Scientific Computing Research (ASCR), the FES community is well positioned to be a major applications area for demonstrating the benefits of next generation exascale computing. For example, the computational FES projects on the ASCR leadership class facility (LCF) supercomputers at Oak Ridge National Laboratory and Argonne National Laboratory are demonstrably among the leading applications in the U.S. computational science community. The FES advances in models, algorithms, and software have and will continue to demonstrate

“applications readiness” of benefit to the ASCR mission to “to develop the algorithms, computer programs and hardware that advance scientific research.” The FES domain can be expected to continue to “show the way” for other application areas by its prominent role in cross-cutting ASCR-led HPC programs such as the aforementioned SciDAC and also the INCITE (Innovative and Novel Computational Impact on Experiment) programs. Of particular current interest to ASCR is the availability of an impressive suite of well-diagnosed FES experimental facilities and associated large foundational data sets. Properly aligned with the FSP, this will help drive UQ (Uncertainty Quantification) research and development that features sensitivity analysis.

5. FSP Approach

The current FSP planning activity has involved the identification and prioritization of those critical scientific challenges (“science drivers”) that will best accelerate progress toward the delivery of fusion power. Key to this process has been the “gaps analysis” of associated software tools with respect to needed improvements in not only the physics representations but also the maturity of the computer science and applied mathematics incorporated. The associated computational challenge is to deliver a suite of predictive integrated FES simulation capabilities that are properly validated against experiments in regimes relevant for producing practical fusion energy and producing a living-scientific-road-map that identifies compelling deliverables needed for progress on the science drivers. These include:

(1) Disruption Avoidance & Mitigation:

Disruptions are large-scale macroscopic events leading to rapid termination of plasma discharges, including severe impulsive heat loads damaging material components. Avoiding or at least mitigating them is critical because ITER can sustain at most a very small number of full current disruptions.

(2) Pedestals:

These regions at the plasma edge are characterized by steep spatial gradients and have a major influence on core confinement and on transient heat loads. Improved understanding is needed for reliable predictions of core confinement and for controlling the frequency and size of edge localized mode (ELM) crashes.

(3) Plasma Boundary Physics:

The plasma boundary layer is the region at the plasma periphery where complex plasma-materials interactions take place. Improved understanding is needed for mitigating heat loads that impact: (i) divertor design and operational strategies; (ii) erosion of divertor and plasma facing components; and (iii) tritium retention

and removal.

(4) Wave-Particle Interactions (Energetic Particles & RF Physics):

These are dynamical interactions between energetic particles and electromagnetic waves in an MFE plasma that impact the efficacy of auxiliary heating and the fast-particle confinement of fusion products (alpha particles at 3.5 MeV) and of supra-thermal particles from radio-frequency (RF) and energetic neutral beam ion heating (NBI). Improved predictive capability is essential for ensuring steady-state (long-pulse) performance in burning plasmas such as ITER.

(5) Core Profiles:

These are plasma profiles in the core confinement region of magnetically confined plasmas. Improved predictive capability for the temperature, density, current, and rotation profiles in the plasma core, including the internal transport barrier region is needed to determine operational limits (e.g., sustainable plasma pressure) and plasma performance (e.g., fusion yield), and also to provide confidence in extrapolating core confinement predictions to future devices.

(6) Whole Device Modeling (WDM):

This is an experimentally validated predictive capability to simulate/model the entire MFE system, including the dynamics of the other 5 previously described MFE issues/challenges. Development of a suite of high-physics fidelity tools/capabilities are essential for enabling reliable scenario/operating design for existing and planned (e.g., ITER) experimental systems as well as for future devices such as a demonstration reactor (DEMO).

With regard to prioritization, the following metrics are followed in choosing the FSP areas of focus:

(1) Need for Multi-scale, Multi-physics Integration: The proposed topic should be outside the areas of concentration in existing modeling programs. Significant advances on (or actual solutions of) the problems posed should demonstrate that the FSP "is more than the sum of its parts."

(2) Importance and urgency: The solution of the problems posed should be integral to the creation of the knowledge base needed for the FES mission leading to "an economically and environmentally attractive fusion energy source. Urgency is related to schedules, dependencies, and critical paths for program elements that the FSP would support

(3) Readiness and Tractability: The underlying physics base — together with the requisite base from applied math and computer science together with access

to the needed computing platforms — should be sufficient to carry out research at the outset of the FSP. Requirements here include the need for the FSP to provide a clear "living roadmap" to track substantive progress on the research topics targeted – with the expectation that new results will impact on-going research at an early date, and

(4) Opportunity for New Lines of Research: The associated R&D for the proposed topic should offer clear opportunities for delivering new insights or potential breakthroughs — particularly those not accessible by other means.

In addition, it is important to ensure "buy-in" from the "customer-base" for FSP products. To do so, the final Integrated Science Application (ISA) plans are expected to explain/highlight which user communities are interested in the FSP software capability proposed for development and with what level of urgency. This implies that the FSP planning documents should reflect a realistic level of "market analysis" of the needs articulated, for example, by the ITPA, the U.S. Fusion Facilities Coordinating Committee (FFCC), and ITER.

Since virtually all of the key FES problems are "interdisciplinary" and highly complex, associated scientific progress demands HPC-enabled integrated modeling. In this sense, a successful FSP will benefit ITER and all areas of FES. In the development of the program definition of the FSP, integrated science applications (ISA) teams targeting the gaps in the "science drivers" areas are coordinating closely with the teams working on the planning of the needed physics components, integration frameworks, and the experimental validation campaigns to ensure that the FSP goals are achieved in a timely and effective manner.

The FSP is conceived to be an ongoing program dedicated to research and development activities for producing experimentally validated computer simulation software to help accelerate progress in magnetic fusion energy research. The initial goals of the FSP have been established as deliverable products or services, and the program's success will be measured in part by progress toward, and eventual delivery of, those products and services. The actual measure of success will be the effectiveness of those products and services once they become operational. As such, the program will be managed as a mixed life-cycle effort whereby some program activities will be managed using project management techniques and some managed as ongoing operations.

Some specific examples of expected advances that are needed to enable a comprehensive integrated modeling capability include:

- Effective coupling of state-of-the-art codes for the

plasma core and the plasma edge region.

- Effective coupling of state-of-the-art codes for MHD dynamics with codes for auxiliary heating of the plasma via RF waves.
- Development of more realistic reduced models based on results obtained from the DNS-type (direct numerical simulation) major codes that use petascale capabilities.
- Development of advanced frameworks and workflow management needed for code coupling.

It is expected that the FSP will continue to actively engage in information exchange targeting potential areas of fruitful collaborations with integrated modeling efforts outside the US. Such activities in the EU and Japan have primarily focused on significantly extending the capabilities of currently operating production models. The FSP planning activity has also made excellent progress in engaging the major U.S. experimental facilities (DIII-D, C-MOD, and NSTX) to help provide validation foundations needed to establish the physics fidelity of theoretical and simulation models. Favorable discussions have also been initiated with international facilities that have capabilities unavailable in the U.S., including: (i) Asia with superconducting long-pulse facilities such as KSTAR in Korea and EAST in China; (ii) the EU with DT experiments in JET; and (iii) with the international database coordinated by ITER.

6. Major Computational Challenges

A dominant trend emerging in the next-generation (“ultrascale”) hardware development area is to continue to add more CPU (central processing unit) cores onto the same chip and also use accelerators such as GPU’s (graphics processing unit) to deliver high aggregate computing performance. Current estimates are that the number of cores per chip is expected to increase by an order of magnitude within five years and two orders of magnitude in a decade. Together with the high-density, low-power packaging approach to construct large-scale parallel computers, this trend indicates that an exascale machine in the next decade could reach 10 million CPU cores. The formidable challenge here is of course to develop new methods to effectively utilize such dramatically increased parallel computing power. This will be necessary to achieve accelerated scientific discovery in fusion energy sciences as well as many other applications domains.

Some specific examples of outstanding challenges in the fusion energy science application area include:

- Efficient scaling of MHD codes beyond terascale levels to enable higher resolution simulations with associated greater physics fidelity
- Efficient extension of global PIC codes into fully electromagnetic regimes to capture the fine-scale dynamics relevant not only to transport but also to help verify the physics fidelity of MHD codes in the long-mean-free-path regimes appropriate for fusion reactors.
- Mastery of data management capabilities to help with the development (including debugging) of advanced integrated codes.
- Development of innovative data analysis and visualization capabilities to deal with increasingly huge amounts of data generated in simulations at the petascale and beyond.

A possible approach to dealing with the many-core challenge in exascale supercomputers is to consider hybrid-type programming models between message passing and shared memory. As the future many-core processor chips are likely to support shared memory, and parallel computers are expected to support a form of message-passing (e.g, MPI, UPC, etc.) among nodes, it will be necessary to consider algorithms and programming techniques to construct parallel programs where the code that runs on an individual node uses the multi-threaded, shared-memory programming model (OpenMP, CUDA, OpenCL, etc.), while also using the message-passing programming model (such as MPI or UPC) to communicate between nodes. The associated transition to many-core programming is characterized by very low latency but limited bandwidth challenges for CPU’s – a key aspect of future exascale co-design considerations in the FES as well as many other application domains.

In order to achieve the desired accelerated progress, co-design investments must be made to develop:

- Compilers to decompose the code running on a single node into fine-grained computation tasks to utilize the collection of cores on a single chip.
 - Highly efficient run-time system to schedule fine-grained tasks to optimize for available concurrencies and to maximize on-chip cache locality to overcome off-chip memory latency and bandwidth constrains.
 - Hybrid-type programming environments for existing MPI and UPC by incorporating the compilation and run-time systems for OpenMP, CUDA, and OpenCL.
- If proper investments in research efforts are made, specific fusion energy science state-of-the-art codes can be transformed to versions using a hybrid programming framework and then systematically exercised to demonstrate and to test the effectiveness of this proposed

paradigm. In the event that such methods for optimally utilizing many-core systems prove to be truly effective, then the expected outcome would be that such codes can proceed to realize their high potential for achieving new scientific discoveries with first multi-petascale and then exascale computing power. Other examples of expected outcomes if dedicated efforts are properly supported include:

- Clear demonstration of the ability to effectively integrate (beginning with two-way coupling) advanced codes to deliver new physics insights;
- Significant progress on the ability to reliably predict the important Edge Localized Modes (ELM) behavior at the plasma periphery; and
- Significant improvement in the physics fidelity of “full device modeling” to enable timely achievement of reliable predictive capability for ITER.

7. Conclusions

In general, reliable whole-device modeling capabilities in Fusion Energy Sciences will surely demand computing resources at the petascale range and beyond to address ITER burning plasma issues [4]. Even more powerful exascale platforms will be needed to meet the future challenges of designing a demonstration fusion reactor (DEMO). With ITER and leadership class computing being two of the most prominent international missions, the development of an experimentally-validated integrated modeling suite of tools targeted by the Fusion Simulation Program (FSP), is a timely exascale-relevant project for producing a modern HPC capability for fusion. If successful, the FSP should prove to be of major benefit to strategic considerations for accelerating progress toward the delivery of magnetic fusion energy for the world.

Acknowledgements

This work was supported by DoE Contract No. DE-AC02-09CH11466.

The author is especially grateful for many productive and stimulating discussions with his colleagues on the U.S. Fusion Simulation Program (FSP) Planning Team [M. Greenwald, A. Siegel, D. Kothe, J. Cary, X. Tang, V. Chan, A. Kritz, J. Hittinger, and others (<http://www.pppl.gov/fsp/>)].

References

- [1] W.M. Tang and V.S. Chan, *Plasma Phys. Control. Fusion* **47**, No. 2, R1-R34 (2005).
- [2] X. Garbet, Y. Idomura, L. Villard, and T.H. Watanabe, *Nucl. Fusion* **50** (2010) 043002 (30pp) (2010).
- [3] D.A. Batchelor, M. Beck, A. Becoulet, R.V. Budny, C.S. Chang, P.H. Diamond, J.Q. Dong, G.Y. Fu, A. Fukuyama, T.S. Hahm, D.E. Keyes, Y. Kishimoto, S. Klasky, L.L. Lao, K. Li, Z. Lin, B. Ludaescher, J. Manickam, N. Nakajima, T. Ozeki, N. Podhorszki, W.M. Tang, M.A. Vouk, R.E. Waltz, S.J. Wang, H. R. Wilson, X.Q. Xu, M. Yagi and F. Zonca, *Plasma Science and Technology*. **9**:312-387, 2007.
- [4] W.M. Tang *et al.*, “Scientific Grand Challenges: Fusion Energy Sciences and the Role of Computing at the Extreme Scale,” DoE-SC Peer-reviewed report on major workshop held March 18-20, 2009, Washington D.C. Final Publication PNNL-19404, 212pp (2010). Production support provided by Pacific Northwest National Laboratory for the U.S. Department of Energy’s Office of Advanced Scientific Computing Research and the Office of Fusion Energy Sciences, Washington, D.C. <http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf>.
- [5] W.W. Lee, *Phys. Fluids* **26**, No. 2, 556-562 (1983).
- [6] W.W. Lee, *J. Comput. Phys.* **72**, 243-269 (1987).
- [7] Z. Lin, T.S. Hahm, W.W. Lee, W.M. Tang, and R.B. White, *Science* **281**, 1835-1837 (1998).
- [8] Z. Lin, T.S. Hahm, W.W. Lee, W.M. Tang, and R.B. White, *Phys. Plasmas* **7**, No. 5, 1857-1862 (2000).
- [9] S. Ethier, W.M. Tang, R. Walkup, and L. Oliker, *IBM Journal of Research and Development* **52** (1-2) 105–115 (2008).
- [10] L. Oliker *et al.*, “Scientific Application Performance on Candidate PetaScale Platforms,” *In Proc. IPDPS’07*, Long Beach, CA, March 24-30, 2007.
- [11] S. Ethier, W.M. Tang, and Z. Lin, *J. Phys.: Conf. Series* **16**, 1-15 (2005).
- [12] L. Oliker *et al.*, “Leading computational methods on scalar and vector HEC platforms,” *In Proc. SC2005*, Seattle, WA, Nov 12-18, 2005.
- [13] L. Oliker, A. Canning, J. Carter, J. Shalf, and S. Ethier, “Scientific Computations on Modern Parallel Vector Systems,” *In Proc. SC2004*, Pittsburgh, PA, Nov 6-12, 2004.

The Princeton Plasma Physics Laboratory is operated
by Princeton University under contract
with the U.S. Department of Energy.

Information Services
Princeton Plasma Physics Laboratory
P.O. Box 451
Princeton, NJ 08543

Phone: 609-243-2245
Fax: 609-243-2751
e-mail: pppl_info@pppl.gov
Internet Address: <http://www.pppl.gov>