PREPARED FOR THE U.S. DEPARTMENT OF ENERGY,
UNDER CONTRACT DE-AC02-76CH03073

PPPL-3879
UC-70

PPPL-3879

Plasma Science Contribution to the SCaLeS Report

by

S.C. Jardin

October 2003

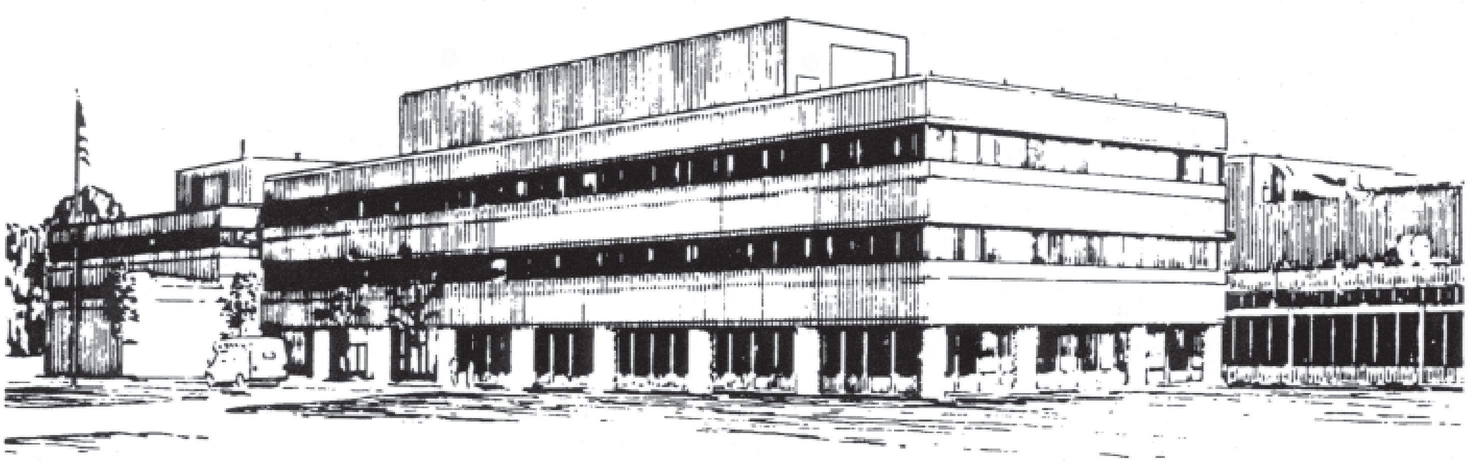

PRINCETON PLASMA PHYSICS LABORATORY
PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY

Availability

This report is posted on the U.S. Department of Energy's Princeton Plasma Physics Laboratory Publications and Reports web site in Fiscal Year 2004. The home page for PPPL Reports and Publications is: http://www.pppl.gov/pub_report/

DOE and DOE Contractors can obtain copies of this report from:

U.S. Department of Energy
Office of Scientific and Technical Information
DOE Technical Information Services (DTIS)
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401

Fax: (865) 576-5728

Email: reports@adonis.osti.gov

This report is available to the general public from:

National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

Telephone: 1-800-553-6847 or
(703) 605-6000

Fax: (703) 321-8547

Internet: http://www.ntis.gov/ordering.htm

Plasma Science Contribution to the SCaLeS Report

S. C. Jardin

Princeton Plasma Physics Laboratory
Princeton, NJ 08543

Abstract

In June of 2003, about 250 computational scientists and mathematicians being funded by the DOE Office of Science met in Arlington, VA to attend a 2-day workshop on the Science Case for Large-scale Simulation (SCaLeS). This document was the output of the Plasma Science Section of that workshop. The conclusion is that exciting and important progress can be made in the field of Plasma Science if computer power continues to grow and algorithmic development continues to occur at the rate that it has in the past. Full simulations of burning plasma experiments could be possible in the 5-10 year time frame if an aggressive growth program is launched in this area.

What is Plasma Science?

Plasmas are very hot gases in which the individual atoms have broken up into a collection of electrons and atomic nuclei (or ions). Sometimes called the fourth state of matter, plasmas exhibit a rich variety of complex, collective phenomena. The sun and the stars are predominantly plasma, as is the composition of over 99% of the visible universe.

Plasmas interact strongly with magnetic fields. These fields can be imposed externally or be created by electrical currents flowing in the plasma itself. A plasma configuration can become unstable if the electrical currents, the temperatures, the densities, or their gradients exceed critical values relative to one another and to the strength and geometry of the magnetic field. As a result of these instabilities, the plasma will rearrange itself to seek a stable configuration where these criteria are satisfied. This rearrangement can occur on a global scale (Magnetohydrodynamic or MHD stability), or on a fine scale (plasma micro-turbulence).

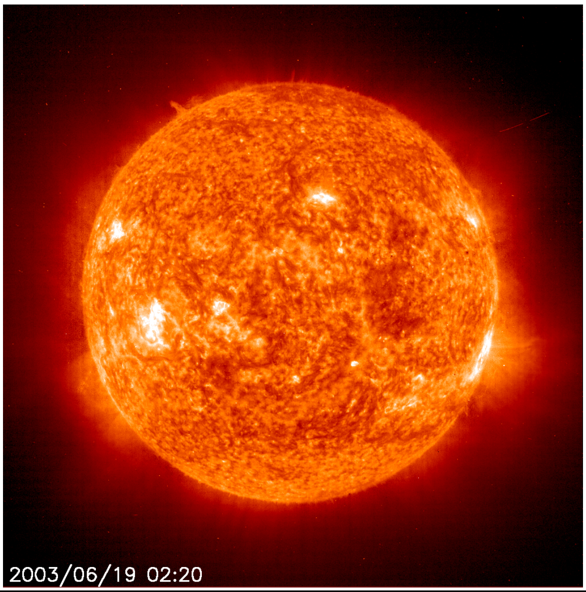

Figure 1: View of the sun from the SOHO satellite. Over 99% of the visible universe is plasma.

Radio-frequency (RF) electromagnetic waves interact strongly with plasmas. They are reflected, absorbed and transmitted in plasma depending on their frequency, the plasma temperature and density, and the orientation of the wave fields relative to the background magnetic field.

Plasma science seeks to understand the physical processes which underlie these phenomena.

Impact on science and society

Although plasmas play an important role in many aspects of everyday life, e.g. neon signs, plasma video displays, spark plugs, and flames, a major focus of research in plasma science is the quest for harnessing fusion energy. The development of a secure and reliable energy system that is environmentally and economically sustainable is one of the most formidable scientific and technological challenges facing the world in the twenty-first century. The vast supplies of deuterium fuel in the oceans and the absence of long-term radiation, CO2 generation, and weapons proliferation concerns makes fusion the preferred choice for meeting the energy needs of future generations.

The DOE Office of Fusion Energy Sciences (OFES) supports an active research program in fusion energy science with three major U.S. Magnetic Fusion Energy (MFE) experiments underway, and is currently negotiating a role for the U.S. in the upcoming ITER burning plasma experiment. The U.S. also supports a large magnetic fusion theory effort which has a long history of being at the cutting-edge of computational physics research. In fact, the present National Energy Research Scientific Computing Center (NERSC) is an outgrowth of the MFE computer center, MFECC, which was established in the late 1970s as the first national supercomputer center.

In MFE experiments, high-temperature (100 million degrees centigrade) plasmas are produced in the laboratory in order to create the conditions where hydrogen isotopes (deuterium and tritium) can undergo nuclear fusion and release energy (the same process that fuels our sun). Devices called tokamaks and stellarators are "magnetic bottles" that confine the hot plasma away from material walls, allowing fusion to occur. Unfortunately, confining the ultra-hot plasma is a daunting technical challenge. The level of micro-turbulence in the plasma determines the amount of time it takes for the plasma to "leak out" of the confinement region. Also, global stability considerations limit the amount of plasma a given magnetic configuration can confine and thus determines the maximum fusion rate and power output.

A complementary approach to MFE is called Inertial Fusion Energy (IFE). DOE's IFE program, also within OFES, is coordinated with, and gains leverage from, the much larger Inertial Confinement Fusion (ICF) program of the National Nuclear Security Administration (NNSA). In IFE, intense beams of particles (ion-beam fusion) or light (laser fusion) are focused on small targets which contain pellets of frozen heavy-hydrogen. When these pellets are imploded sufficiently rapidly and symmetrically, the conditions for a small nuclear fusion "explosion" are created. These explosions release substantial energy, but are small enough that their energy can be confined within the fusion chamber, where it can be converted to a useful form. Plasma physics issues arise in the beam itself in the case of ion-beam fusion, in obtaining high compression ratios and maintaining symmetry in the target, and in an advanced concept known as fast-ignition.

Plasma science is also of great importance in understanding crucial interactions between the sun and the Earth. Plasma is always being emitted from the sun in the form of a supersonic wind called the "solar wind". In addition to the solar wind, plasma in the sun's outer atmosphere, called the "corona", can undergo sudden and violent activity in the form of "coronal mass ejections" and "solar flares", examples of which can be seen in Figure 1. As a result of these activities, billions of tons of matter and intense energetic particles can be thrown out of the solar corona into outer space, causing "storms" that can disturb significantly the near-Earth environment. All of the various phenomenon that occur in the near-Earth environment whose behavior and interactions directly affect the planet and human technologies on and in orbit around it make up "space weather". Space weather can have significant effects for several Earth-based technologies such as satellites, communications and navigation systems, and radiation exposure in manned space missions.

Scientific opportunities

Computational modeling currently plays an essential role in all aspects of plasma physics research. Perhaps nowhere is this as evident as it is in magnetic fusion energy (MFE) research where simulation models are actively being improved, tested and applied to the interpretation of data and to the design of new experiments. Improvements in the modeling comes in the form of both more complete models that include better descriptions of the physical processes and more efficient models that use advanced algorithms.

Present capability is such that we can apply our most complete computational models to realistically simulate both nonlinear macroscopic stability and microscopic turbulent transport in the smaller fusion experiments that exist today, at least for short times. Anticipated increases in both hardware and algorithms during the next 5-10+ years will enable application of even more advanced models to the largest present-day experiments and to the proposed burning plasma experiments such as ITER. (See Figure 2 and the discussion below).

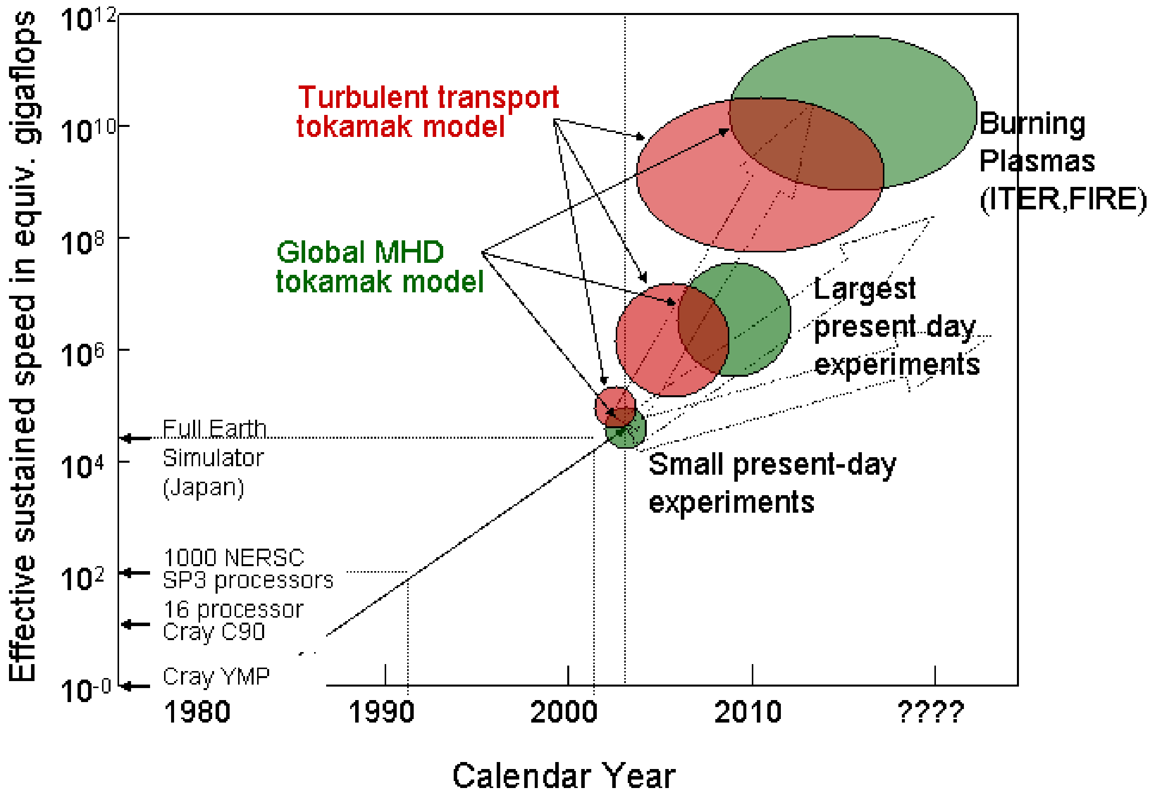

Figure 2: Factors of 100-10,000 in effective sustained speed are required to do complete modeling of proposed MFE burning plasma experiments.

A number of advances in the formulation and algorithms have complemented the increases in hardware speeds to provide vastly improved capability today than what was possible 30 years ago (see Figure 3). We expect this trend to continue into the future. This rate of increase of effective capability is essential to meet the anticipated modeling demands of fusion energy research as described below.

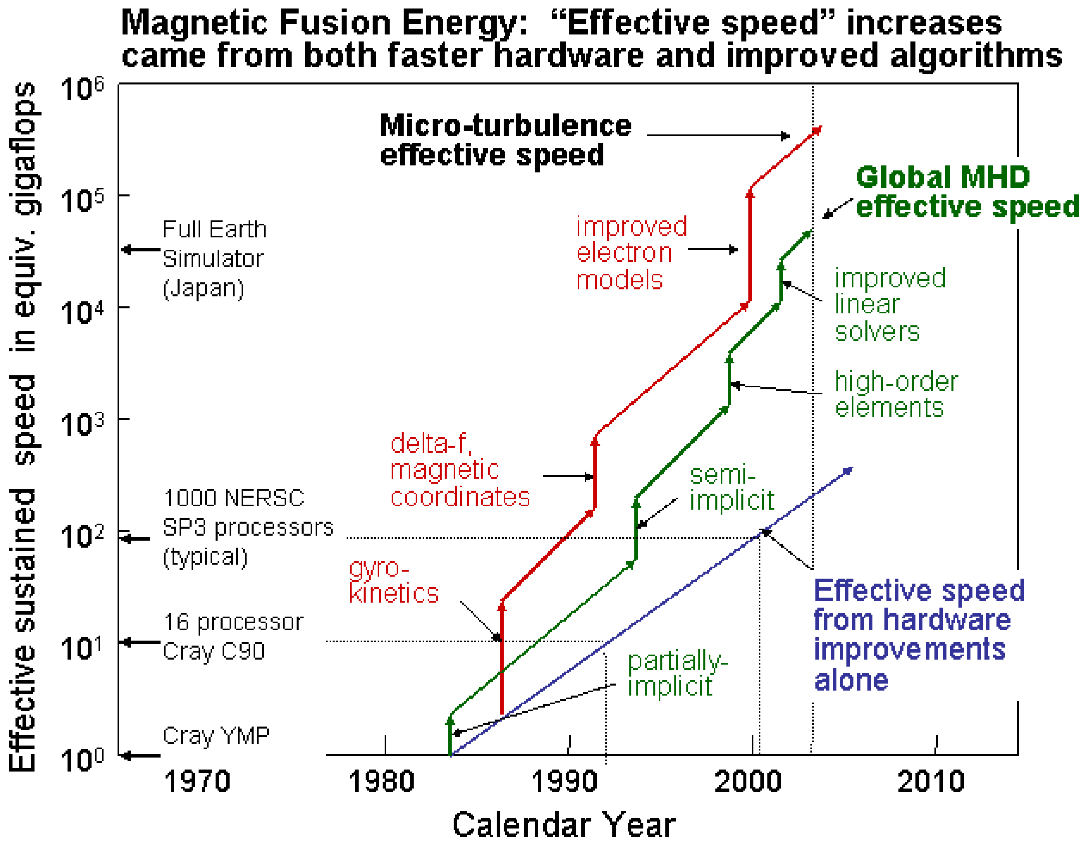

Figure 3: Magnetic Fusion Energy: "Effective speed" increases came from both faster hardware and improved algorithms.

The present thrust in computational plasma science is to merge together the now separate macroscopic and microscopic models, and to extend the physical realism of these by the inclusion of detailed models of such phenomena as RF heating and atomic and molecular physical processes (important in plasma-wall interactions), so as to provide a true integrated computational model of a fusion experiment. Such an integrated modeling capability will greatly facilitate the process whereby plasma scientists develop understanding and insights into these amazingly complex systems that will be critical in realizing the long term goal of creating an environmentally and economically sustainable source of energy.

A number of external drivers are at work to make this time an especially opportune one for accelerating our capabilities in computational modeling of plasma. In MFE, the international ITER experiment is scheduled to begin its 10-year construction phase in 2006. There is a clear opportunity for the U.S. to take the lead in the computational modeling of this device, putting the U.S. in a strong position to influence the choice of diagnostic hardware installed and the operational planning of the experiments, and to take a lead in the subsequent phase of data interpretation. Furthermore, a comprehensive simulation model such as envisioned in the Fusion Simulation Project is felt to be essential in developing a demonstration fusion power plant, to follow ITER, by effectively synthesizing results obtained in ITER with those from other non-burning experiments which will be evaluating alternate MFE configurations during this same time period.

The IFE community expects that favorable results from the National Ignition Facility (NIF) in the next few years will further validate the models used in their target designs and will give an extra impetus to proceed with an Integrated Research Experiment in the 2015 time-frame. The space-weather community is anticipating an unprecedented influx of high-quality data from the NASA Magnetospheric Multi-scale Missions in 2009 and need to have computational predictions to be able to compare with these measurements.

Research issues

The plasmas in modern magnetic fusion experiments are typically not quiescent, but exhibit macroscopic motions that can affect their performance, and in some cases can lead to catastrophic collapse of the discharge. (See Figure 4)

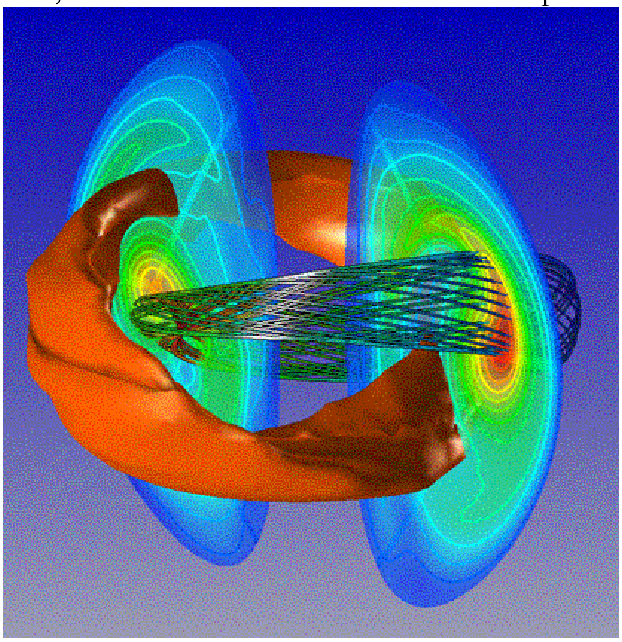

Figure 4: Extended-MHD calculation of an internal mode in NSTX (courtesy W. Park)

The modeling of such dynamics for realistic experimental parameters requires an integration of fluid and kinetic physics in a complex magnetic geometry as described by the extended-MHD equations. The magnetic field, required for confinement, imposes a large anisotropy to the problem. However, the key challenge in performing computations relevant to the hot plasmas of modern fusion experiments is to increase the dimensionless parameter characterizing inverse plasma collisionality, the Lundquist number, S. Present global MHD calculations are limited to Lundquist numbers $S < 10^7$ and problem times $T < 1$ msec. These values are adequate for modeling small low-temperature experiments, but are several orders of magnitude less than what are required to accurately simulate the largest of the existing fusion experiments. Several more orders of magnitude would be required to simulate “ITER-class” burning plasma experiments.

The confinement of energy and particles in fusion plasmas is often significantly degraded by turbulence associated with small spatial-scale plasma instabilities driven by gradients in the plasma pressure. (See Figure 5)

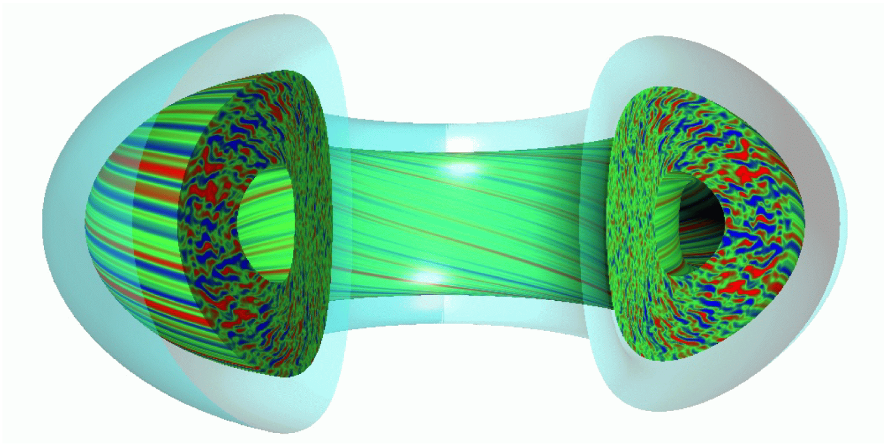

Figure 5: Electric potential structure associated with micro-turbulence in tokamaks.

The detailed physics of the growth and saturation of these instabilities, their impact on plasma confinement, and the development of an understanding of how such turbulence might be controlled remain unsolved problems for which we have only glimpses of understanding. At the present time roughly 10^{-3} s of a turbulent discharge can be modeled. This time needs to be increased by a factor of 10-100 to address relevant time scales in the largest experiments and even more for ITER-class burning plasma experiments.

The bulk of the plasma turbulence results today have been obtained with an accurate model of ion dynamics (kinetic ions), which play a dominant role, but with a simplified (adiabatic or fluid) model for the electrons. This is not adequate for making quantitative predictions for real experiments. Early simulations of electron and electromagnetic effects reveal important dynamics on smaller and faster scales than what are encountered in electrostatic calculations. Simulations which can simultaneously resolve ion, electron, and electromagnetic-scale interactions necessitate an increase in computing resources of 50-100.

The scientific issues that arise when modeling a magnetic fusion experiment encompass a wide range of disciplines including those mentioned above, as well as others. However the dynamics of high temperature plasma does not respect these categorizations and an understanding of overall plasma performance requires combining all of these disciplines in an integrated simulation that includes interactions between phenomena which were previously studied as essentially separate disciplinary problems. To achieve the ultimate goal of such an integrated approach, we must simulate the evolution of the 3D distribution of the plasma temperature, density, current and magnetic field on long time-scales in a way that includes all the relevant physical processes active at the shorter time-scales. While this is a long-term and ambitious goal, the program now stands ready to begin such cross-disciplinary studies and to increase the physics content of existing integrated codes. To accelerate this process the fusion community is engaged in a study laying the groundwork for a major initiative referred to as the Fusion Simulation Project. [See sidebar] A major requirement for this endeavor is access to significantly increased computing power.

Progress in all key physics areas of Inertial Fusion Energy (IFE), including the "drivers" which impart the energy to the fusion fuel, the targets, and also the fast-ignition concept could be dramatically accelerated by increased computing resources. The principal IFE driver approach supported in the Office of Science consists of beams of heavy ions produced by linear induction accelerators. These intense beams are non-neutral plasmas that exhibit collective behaviors dominated by space-charge forces; demanding a self-consistent, integrated treatment from the source to the target.

Similarly, 3-D simulations of targets are required. Finally, there is broad international interest in fast ignition, which uses a separate short-pulse laser to ignite the compressed fuel, reducing the total required input energy.

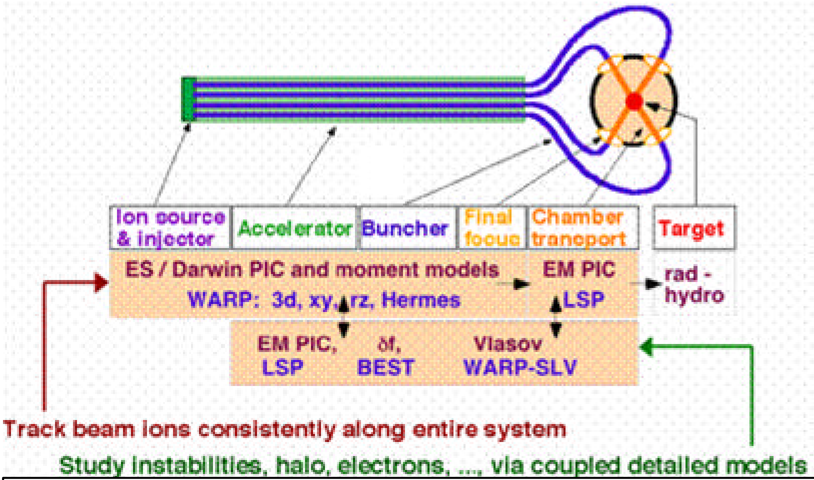

Figure 6: Heavy Ion driver simulation strategy.

Space weather simulations typically couple physical processes on the very large solar-terrestrial scales to small scales that are one-thousandth of the Earth's radius. It is necessary to carry out such calculations for several hours of real time in order to be able to predict even short-term space weather. (Such a simulation would require of the order of 10^{21} Flops and run for several months at close to peak performance on the 40 TFlop Japanese Earth Simulator.) Furthermore, as in MFE research, our physical understanding of plasma macro- and micro-instabilities and their implications for plasma stability and transport remain poorly understood, and the interplay of these effects in a complex integrated model can be done only by means of computer simulations.

Technology barriers

The existing large-scale plasma science codes, both kinetic and (extended-MHD) fluid-based codes, typically exhibit relatively low ratios of sustained to peak performance (2% to 10%) on the current generation of IBM SP-type machines. They are limited by memory bandwidth and latency rather than by raw processor speed. The PIC kinetic codes would also benefit from computers with hardware support for gather-scatter and scatter-add operations. The codes typically do not scale well above 1000 processors (for strong scaling) on existing architectures, primarily due to latency in the interprocessor interconnects.

There are many algorithmic improvements that could still be made to the plasma science codes. The fluid codes are dominated by sparse-matrix operations, and improvements here would be of immediate benefit. There is a need for application-specific pre-conditioners for strong anisotropy and multi-scale phenomena. Initial results with high-order and spectral finite elements look very promising for efficiently representing multi-spatial scales and strong anisotropy. Non-linear implicit techniques show promise for dealing with the multi-temporal scales. The hybrid particle/fluid description is a particularly efficient representation for many applications where a small component of high-energy particles affects the global stability properties of the plasma. This should be developed further. Advanced adaptive mesh refinement (AMR) algorithms also show promise but must be extended to handle implicit equations and hybrid calculations.

The large amounts of data generated by plasma science simulations are inconsistent with the traditional data-management and visualization tools now in use. Tools for large-scale data and meta-data management and to facilitate remote visualization are needed. Also needed are efficient means for check pointing and restarting large jobs and software tools that facilitate managing the complexity of large distributed software projects.

Resources required

An increase of 100-1000 in computing power would lead to a step-wise enhancement in each of the sub-fields mentioned above, enabling dramatic new capabilities that can move plasma science to the next level. The increase should take many forms, from desktop to flag-ship facility. Increased compute power is not only needed for increased resolution. It is also needed to enable integrating model components, increasing model fidelity, executing longer runs, increasing the number of runs in an ensemble to improve statistics or to investigate multi-dimensional parameter spaces, and increasing the number of applications of a given capability. We plot in Figure 7 the approximate ratios of compute power (in terms of TeraFLOPs for an entire calculation) versus memory requirement in Gigabytes for typical existing and anticipated calculations in Plasma Science.

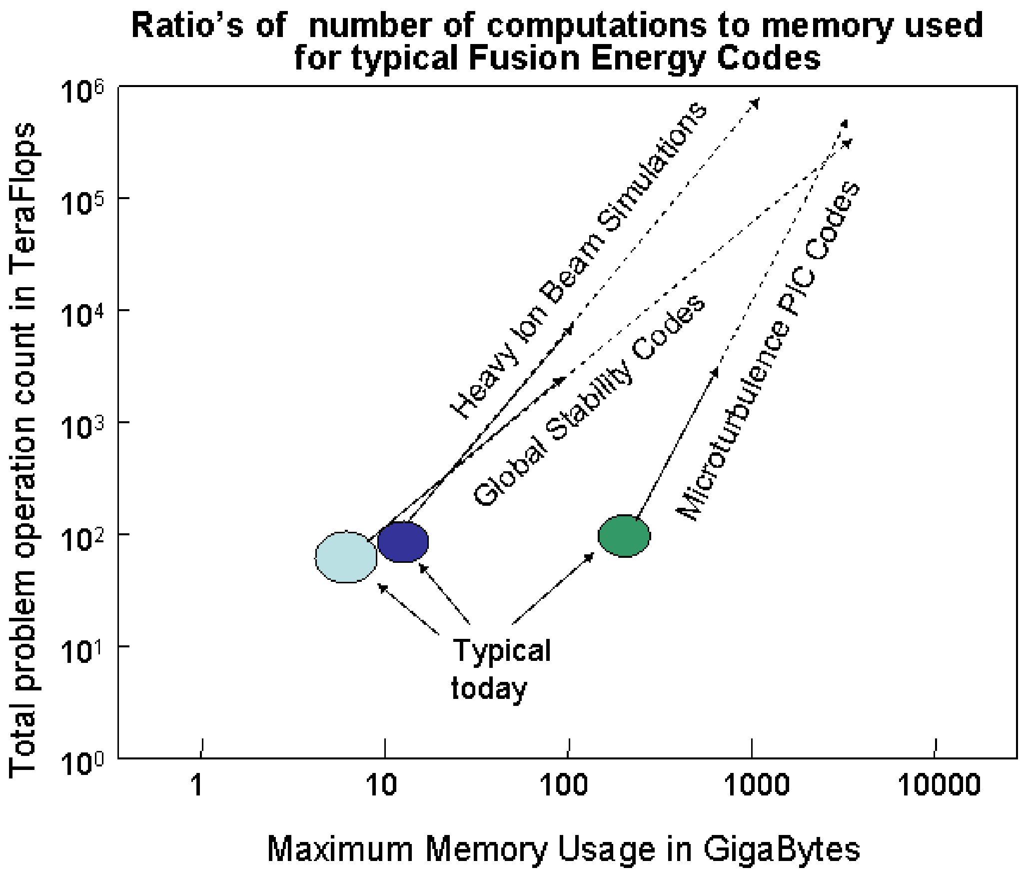

Figure 7: Total flops/calculation and total memory usage for typical existing and anticipated calculations in plasma science.

This should serve as a guide to the type of large-scale computing hardware required. To be effective, this hardware increase must also be accompanied by an increase in human resources in the form of teams consisting of computational physicists, applied mathematicians and computer scientists.

Future capabilities and MFE community computing requirements

	Today	5-years	10-years
Target capability	• resistive MHD, small experiments, short times • turbulence in core of small experiments, • turbulence in core of ITER for short times • fluid edge turbulence in existing experiments	• resistive MHD in DIII-D • 2-fluid MHD in small experiments • routine turbulence simulation of existing experiments • turbulence in ITER core • kinetic edge turbulence in existing experiments	• resistive MHD in ITER • 2-fluid MHD in DIII-D • turbulence simulation of ITER, core to edge, including transport barriers, short times
Required flops	5×10^{18}	5×10^{19}	5×10^{20}
Aggressive capability		• resistive MHD in ITER • 2-fluid MHD in DIII-D • turbulence simulation of ITER, core to edge, including transport barriers, short times	• 2-fluid MHD in ITER • turbulence in ITER, core to edge, long times • unify microscopic and macroscopic scales
Required flops		5×10^{20}	5×10^{22}
Minimal capability		• resistive MHD, small experiments, longer times • turbulence, small experiments, short times	• resistive MHD in DIII-D • 2-fluid MHD in small experiments • routine turbulence simulation of existing experiments • turbulence in ITER core • kinetic edge turbulence in existing experiments
Required flops		10^{19}	5×10^{19}

Note: flops = total (actual) floating point operations in 1 year by entire community

Figure 8: Resources required in Total flops/year for minimal, target, and aggressive capability in MFE computational physics research.

In MFE, for example, we anticipate that an increase of 100-1000 in computing power will enable fusion researchers to quantitatively predict the onset conditions, strength, and nonlinear saturation mechanisms of both micro-scale and macro-scale instabilities, a major step in understanding how to control them. This same level of increase in computing power would allow researchers to begin developing fully integrated simulations of fusion systems that span the scales from micro to macro. Such an integrated simulation capability would dramatically enhance the utilization of a burning fusion device in particular and the optimization of fusion energy development in general, and would serve as an intellectual integrator of the broad range of physics phenomena occurring in advanced tokamaks.

In space weather, with a factor of 100-1000 increase in current computational speed, it would become feasible to carry out global space weather simulations that couple large solar-terrestrial scales to much smaller scales involving ion dynamics at Lundquist numbers $S < 10^6$. This cannot be done at the present time with existing supercomputers in the US, and yet is necessary if space weather models aspire to capture accurately the physics of collisionless space plasmas.

Besides the hardware costs, it is anticipated that a major new computing initiative directed at developing an integrated computational model of a MFE fusion reactor would require an influx of new funding of about $20 M/year for 5-15 years to support an integrated team of physicists, mathematicians, and computer scientists to develop the large scientific application software required to provide high fidelity simulations of the reactor. Similar initiatives in IFE would require $5 M/year, and in space-weather $3 M/year.

Metrics of success

The first metric of success of this simulation effort can be measured by the degree to which simulation results agree with existing and anticipated high-resolution experimental measurements. There are a wide range of experimental results for comparison that are available now or will be available in the near future. New MFE imaging diagnostics will facilitate comparisons with turbulence and RF calculations. The proposed integrated beam experiment will validate IFE codes. Satellite measurements will be compared with the space-weather calculations. Successes in reproducing these measurements from simulated results will be a dramatic validation of this field's capability and of computational physics in general.

However the true success of this endeavor will be measured by the degree that these simulation codes are relied upon for optimization of experimental operations and the design and optimization of the next generation of experiments and communication systems. In fusion energy, these would be the "Demo-class" power-producing fusion reactors. It is expected that the design capability offered by these simulation codes will give the U.S. a significant competitive edge in the design and manufacture of commercial fusion energy power-plants.

Acknowledgments:

The author gratefully acknowledges contributions from his colleagues: Alex Friedman, Amitava Bhattacharjee, Don Batchelor, Bruce Cohen, Bill Dorland, W.W. Lee, Zhihong Lin, Bill Nevins, Doug Post, Hong Qin, Carl Sovinec, and Scott Parker. Special acknowledgements go to Thom Dunning for making several editorial suggestions. This work was supported by United States Department of Energy Contracts No. DE-AC02-76-CHO-3073.

External Distribution

Plasma Research Laboratory, Australian National University, Australia
Professor I.R. Jones, Flinders University, Australia
Professor João Canalle, Instituto de Fisica DEQ/IF - UERJ, Brazil
Mr. Gerson O. Ludwig, Instituto Nacional de Pesquisas, Brazil
Dr. P.H. Sakanaka, Instituto Fisica, Brazil
The Librarian, Culham Laboratory, England
Mrs. S.A. Hutchinson, JET Library, England
Professor M.N. Bussac, Ecole Polytechnique, France
Librarian, Max-Planck-Institut für Plasmaphysik, Germany
Jolan Moldvai, Reports Library, Hungarian Academy of Sciences, Central Research Institute for Physics, Hungary
Dr. P. Kaw, Institute for Plasma Research, India
Ms. P.J. Pathak, Librarian, Institute for Plasma Research, India
Ms. Clelia De Palo, Associazione EURATOM-ENEA, Italy
Dr. G. Grosso, Instituto di Fisica del Plasma, Italy
Librarian, Naka Fusion Research Establishment, JAERI, Japan
Library, Laboratory for Complex Energy Processes, Institute for Advanced Study, Kyoto University, Japan
Research Information Center, National Institute for Fusion Science, Japan
Dr. O. Mitarai, Kyushu Tokai University, Japan
Dr. Jiangang Li, Institute of Plasma Physics, Chinese Academy of Sciences, People's Republic of China
Professor Yuping Huo, School of Physical Science and Technology, People's Republic of China
Library, Academia Sinica, Institute of Plasma Physics, People's Republic of China
Librarian, Institute of Physics, Chinese Academy of Sciences, People's Republic of China
Dr. S. Mirnov, TRINITI, Troitsk, Russian Federation, Russia
Dr. V.S. Strelkov, Kurchatov Institute, Russian Federation, Russia
Professor Peter Lukac, Katedra Fyziky Plazmy MFF UK, Mlynska dolina F-2, Komenskeho Univerzita, SK-842 15 Bratislava, Slovakia
Dr. G.S. Lee, Korea Basic Science Institute, South Korea
Institute for Plasma Research, University of Maryland, USA
Librarian, Fusion Energy Division, Oak Ridge National Laboratory, USA
Librarian, Institute of Fusion Studies, University of Texas, USA
Librarian, Magnetic Fusion Program, Lawrence Livermore National Laboratory, USA
Library, General Atomics, USA
Plasma Physics Group, Fusion Energy Research Program, University of California at San Diego, USA
Plasma Physics Library, Columbia University, USA
Alkesh Punjabi, Center for Fusion Research and Training, Hampton University, USA
Dr. W.M. Stacey, Fusion Research Center, Georgia Institute of Technology, USA
Dr. John Willis, U.S. Department of Energy, Office of Fusion Energy Sciences, USA
Mr. Paul H. Wright, Indianapolis, Indiana, USA

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: 609-243-2750
Fax: 609-243-2751
e-mail: pppl_info@pppl.gov
Internet Address: http://www.pppl.gov