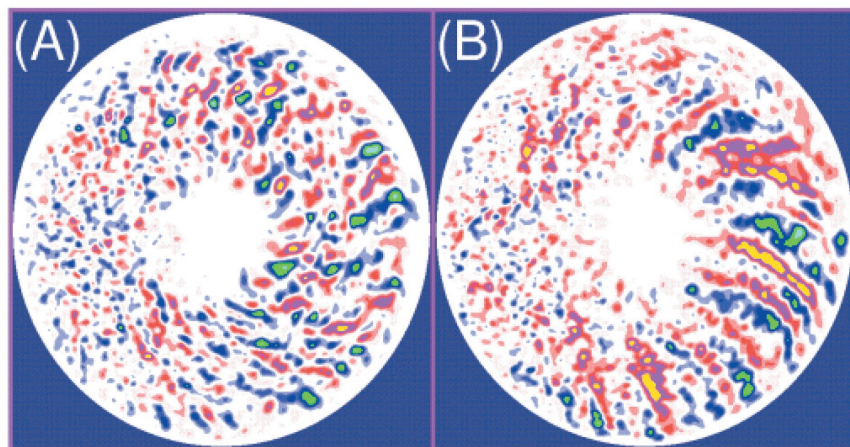


SCIENTIFIC SIMULATION INITIATIVE FOR FUSION ENERGY SCIENCE

3-D Plasma Simulation on New MPPs



With Flow

Without Flow

Turbulence reduction via sheared plasma flow (A), compared to case with flow suppressed (B). Results obtained using full MPP capabilities of CRAY T3E Supercomputer at NERSC [Z. Lin et al., *Science* **281**, 1835 (1998)].

SCIENTIFIC SIMULATION INITIATIVE FOR FUSION ENERGY SCIENCE

Executive Summary

Fusion is the power source of the sun and other stars. It occurs when forms of the lightest atom, hydrogen, combine to make helium in a very hot (100 million degrees centigrade) ionized gas or “plasma.” In this process a small amount of the matter involved in the reaction is converted to a large amount of energy. On earth, fusion could provide a safe, environmentally benign and affordable long-term energy source. Recognizing the potential of fusion energy, the President’s Committee of Advisors on Science and Technology (PCAST) designated fusion as an integral component of the nation’s long-term research portfolio for national energy security and climate change remediation.

Fusion research in the U.S. will be revolutionized by greatly enhanced simulation and modeling capabilities made accessible by terascale computing. Since effective prediction of the properties of energy-producing fusion plasma systems depends on the integration of many complex phenomena spanning vast ranges of time and space scales, advanced scientific computing in tandem with theory and experiment is a vital new tool for discovery. This will enable optimum use of innovative national experiments and support participation in large-scale international facilities. The expected acceleration in scientific understanding and innovation will help provide the scientific foundations needed for key decisions in the 2003–2004 time-frame on major new facilities in this area of research.

The scientific community’s recognition of the importance of a strong U.S. computational effort in fusion plasma research is evident in the recent report from the DOE/NSF Workshop on Advanced Scientific Computing held at the National Academy of Sciences. Building upon a prominent history of advanced computation, which can be traced back to the establishment of the predecessor to the National Energy Research Scientific Computing Center (NERSC) nearly 25 years ago, the U.S. fusion community has been rewarded by impressive advances in the simulation and modeling of plasma behavior. A good example is the effective use of the full power of the most advanced civilian supercomputer at NERSC to produce 3-dimensional particle simulations of turbulence suppression, which is illustrated in the figure on the cover. These calculations, which typically utilized 400 million particles, showed full scalability for 512 processors, indicating a clear path forward to efficient future use of more powerful parallel computers. Such advanced simulations, which help identify and interpret key mechanisms for improving transport in experiments, can aid in revolutionizing our approach to fusion research. Critical insights stimulating greater understanding and innovation will be made accessible by simulations using more advanced physics models covering larger domains for longer periods of time. The combination of terascale computers and increased personnel to develop the requisite computer science and enabling technology capabilities will maximize the return on investment from present facilities and will enable the confident design of the next major experimental steps toward fusion energy. Advanced computing will also stimulate progress in non-fusion plasma applications, including the manufacture of microelectronic components, plasma display panels, and the use of plasmas for waste remediation. Additionally, realistic models will help accelerate the pace of breakthroughs in plasma science applications to energetic beams, to magnetospheric and solar physics, and to astrophysics.

The advanced computational resources targeted by the present initiative can transform fusion research in the next decade by improving scientific understanding of experimental data, by stimulating new theories, and by providing better designs for future facilities. This will lead to a more rapid and cost-effective development of the best approaches to fusion energy production; i.e., a “faster” pace toward “smarter” concepts for more practical fusion systems.

1. Introduction

The development of a secure and reliable energy system that is environmentally and economically sustainable is a truly formidable scientific and technological challenge facing the world in the twenty-first century. In order to prosper in this next century, the United States must remain a leader in research and development on energy supplies and use. Recognizing the potential of fusion energy, the President's Committee of Advisors on Science and Technology recommended strengthening fusion research as a key component of the nation's strategy for energy security and climate change remediation [1]. The vision for the Fusion Scientific Simulation Initiative (SSI) described in this white paper is to catalyze a dramatic change in the way research is presently carried out by enabling the accelerated development of computational tools and techniques that would revolutionize the understanding and design of a fusion energy source. This is made possible by the exciting advances in high performance computing which will allow simulations of more complex phenomena with higher fidelity. Viewed in this way, advanced scientific computing, in tandem with theory and experiment, can be an increasingly powerful new tool for discovery.

Fusion is the power source of the sun and other stars. It occurs when forms of the lightest atom, hydrogen, combine to make helium in a very hot (100 million degrees centigrade) ionized gas or "plasma." Plasmas comprise over 99% of the visible universe and are rich in complex, collective phenomena. In addition to the plasma sciences, research into complex and collective processes are prevalent throughout the broad areas of chemistry, materials sciences, geosciences, and biosciences. For example, in the recent planning document for the Department of Energy Research and Development Portfolio for the Office of Science, plasma physics is included whenever the subject of complex and collective phenomena is mentioned. Understanding the dynamics of processes such as turbulent transport, magnetic reconnection, dynamos, and plasma/material interactions is important for plasmas as well as for many other physical systems. Investigations of plasma turbulence share many common challenges with computational fluid dynamics, and advances in plasma/material interactions will benefit plasma-aided manufacturing. Also, magnetic reconnection, which is of great interest in space and astrophysics studies, and dynamos, which are important in the core of virtually all planetary bodies and stars, are key magnetic confinement research topics. In general, the potential for enhancing the breadth and depth of knowledge via significant cross-cutting alliances with other areas of research is significant.

In thermonuclear plasmas, a small amount of the matter involved in the fusion reaction is converted into a large amount of energy. To produce practical amounts of fusion power in a laboratory environment, a plasma composed of large numbers of fuel nuclei must be heated to temperatures high enough to initiate fusion reactions and be confined well enough to produce net energy. There are two fundamentally different approaches to fusion: magnetic fusion energy (MFE) and inertial fusion energy (IFE). In magnetic fusion, this involves confining a plasma fuel for a sufficient period of time with a magnetic field and heating it to fusion temperatures by injecting high power atomic beams or by radio-frequency waves. In inertial fusion, a continuous sequence of targets containing fusion fuel are heated and compressed to very high densities by high-power lasers or heavy ion beams. A major goal for both of these approaches is to efficiently capture the energy of the products (*e.g.*, neutrons and alpha particles from deuterium-tritium fusion reactions) such that the energy output significantly exceeds the input. This can then be readily used to generate the steam needed to drive the turbines producing electrical power.

Although the fundamental laws that determine the behavior of fusion plasmas, such as Maxwell's equations and those of classical statistical mechanics, are well known, obtaining their solution under realistic conditions is a scientific problem of extraordinary complexity. This is due in large part to the fact that there is an enormous range of temporal and spatial scales involved in fusion plasmas. As illustrated in Fig. 1, the relevant physical processes can span over ten decades in time and space. Effective prediction of the properties of energy-producing fusion plasma systems depends on the successful integration of many complex phenomena spanning these vast ranges. This is a formidable challenge that can only be met with advanced scientific computing in tandem with theory and experiment.

Fusion plasmas are naturally subjected to both large and small-scale disturbances (“instabilities”) which relax the system to a lower energy state. In order to maximize the insulation properties of the magnetic fields, it is necessary to first understand these complex, collective phenomena, and then devise the means to control them. The larger-scale (“macro”) instabilities in magnetically-confined plasmas can produce rapid topological changes in the confining magnetic field resulting in a catastrophic loss of fusion power density. In addition, smaller-scale (“micro”) instabilities can also prevent efficient hot plasma confinement by causing the turbulent transport of energy and particles. Because of the complexity of the kinetic, electromagnetic, and atomic physics equations describing the behavior of fusion plasmas, researchers in both MFE and IFE have a long history of productive use of advanced computation and modeling.

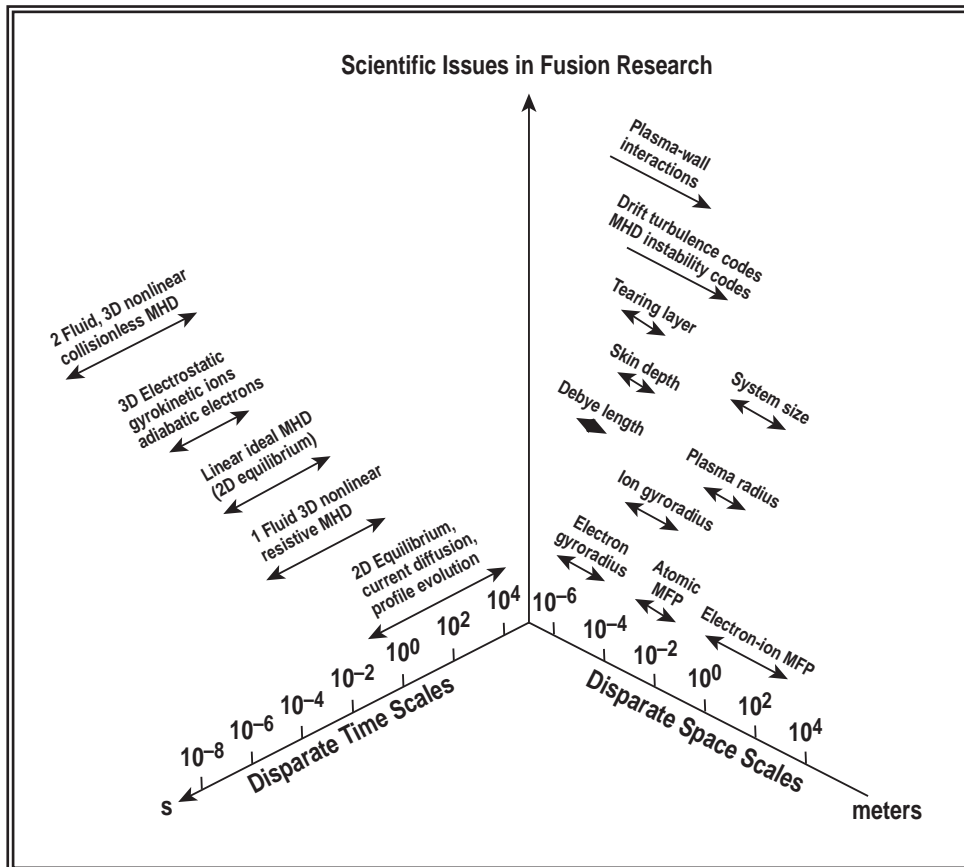


Fig. 1. Disparate scales in fusion research.

The scientific community's recognition of the importance of a strong U.S. computational effort in fusion energy research is evident in the recent report from the DOE/NSF Workshop on Advanced Scientific Computing held at the National Academy of Sciences. Building upon a prominent history of advanced computation, which can be traced back to the establishment of the predecessor to the National Energy Research Scientific Computing Center (NERSC) nearly 25 years ago, the U.S. fusion community has been rewarded by impressive advances in the simulation and modeling of plasma confinement and the interactions of plasma with its surroundings. This has led to major progress in the modeling of turbulence-driven energy transport and macroscopic dynamics of magnetically confined plasmas; and also to major improvements in the simulations of space-charge dominated ion beams for IFE drivers. Some examples of recent successes include simulations of: (1) large-scale disruptions in a fusion-energy-producing toroidal plasma; (2) turbulence suppression with self-generated flows in localized regions of the plasma; and (3) three-dimensional modeling of heavy ion accelerators.

Figure 2 depicts the nonlinear macroscopic simulation of a disruption event which prevented achievement of higher fusion power output in the TFTR tokamak. Insights gained from these studies helped motivate subsequent experiments which better avoided such events. The figures on the cover highlight recent results [2] from advanced particle-in-cell three-dimensional (3D) simulations of the beneficial effects of turbulence suppression. Here self-generated localized ("zonal") plasma flows produce a shearing action which destroys the fluctuating finger-like density contours that promote increased thermal transport. This is also a good example of the effective use of the full power of the most advanced civilian supercomputer at NERSC. These calculations, which typically utilized 400 million particles, showed full scalability for 512 processors, thereby indicating a clear path forward to effective future use of more powerful parallel computers. Advanced simulations of this kind, which help identify and interpret key mechanisms for suppressing turbulence and the associated transport in experiments, can help revolutionize our approach to plasma containment in the fusion energy sciences program. In our third example, 3D particle-in-cell simulation methods have been applied to studies of a heavy-ion based inertial fusion energy system. Valuable information on space-charged-dominated ion beam behavior has been obtained for sections of an IFE driver (Fig. 3). A major future challenge will be to extend such simulations from the source to the target and to then utilize the new capability for designing




Fig. 2. Simulation of high plasma pressure induced disruption in the tokamak TFTR.

a compelling next-step facility. For both of the particle-in-cell simulation examples, valuable new insights into how plasma confinement can be significantly improved will be made accessible by future simulations using more advanced physics models requiring many more particles in larger plasmas for longer periods of time.

The next few years will be a period of important opportunities for both the magnetic and inertial fusion energy programs. For the past decade, the U.S. has been a partner in an international effort to design a magnetically confined burning plasma experiment. During the next 2–3 years, the focus in this area will be to develop lower cost and more attractive

advanced options. A key challenge here is to effectively harvest the most useful scientific discoveries from confinement studies and utilize them in fusion energy producing experiments. The physics guidance made accessible by the fusion SSI is essential for such innovative advances. Operation of the National Spherical Torus Experiment (NSTX) and several smaller innovative magnetic confinement concepts will begin in 1999-2000, and the National Ignition Facility

(NIF) is expected to be online in 2003. The NSTX will be a proof-of-principle experiment to explore the prospects for significant improvement in magnetic confinement, and the NIF will provide definitive data on inertial fusion target physics. The proposed fusion SSI will foster more rapid exploration and understanding of NSTX and other innovative magnetic fusion concepts and will complement the program on NIF by providing enhanced simulation and modeling of IFE systems, such as heavy ion drivers. For each concept there are critical scientific issues that must be solved to achieve a practical fusion system. Effective prediction of the properties of plasma systems depends on the integration of many complex physics phenomena that cannot be deduced from empirical scaling and extrapolation alone. A greatly enhanced theory and modeling effort is essential to extract key physics insights and guide current experiments, to design innovative new devices, and to leverage large-scale international facilities. Despite smaller current investments in research facilities, the U.S. can retain its intellectual leadership in the world fusion program by seizing the opportunity to accelerate the development of realistic simulations of fusion systems. A greatly enhanced computational effort, benchmarked against the best experimental results from national and international experiments, will help provide the scientific and technical foundations needed for key decisions in the 2003–2004 time-frame on major new facilities for fusion research.

The scientific foundations for the Fusion SSI are very strong, and the accelerated development of a realistic predictive simulation capability for both MFE and IFE is now enabled by the availability of state-of-the-art high performance computing platforms. As depicted on Fig. 4, there are two strategic elements in this initiative: direct simulation and integrated modeling. Accelerating the development of the best direct simulation models will require increasing both the sophistication of the mathematical-physics formulation and their resolution. These models fall into two categories: those addressing the microscopic scales which determine plasma transport and those addressing the macroscopic scales associated with rapid large-scale plasma instabilities. New computing resources will allow each category to extend both their simulation domains and the incorporation of more detailed physical effects, ultimately achieving overlapping regions of validity where phenomena, such as macroscopic turbulence and dynamos, can be




Fig. 3. A particle-in-cell simulation used to design a prototype “electrostatic-quadrupole” heavy-ion beam injector at LBNL. Color of each simulation particle denotes the local potential, relative to the on-axis value.

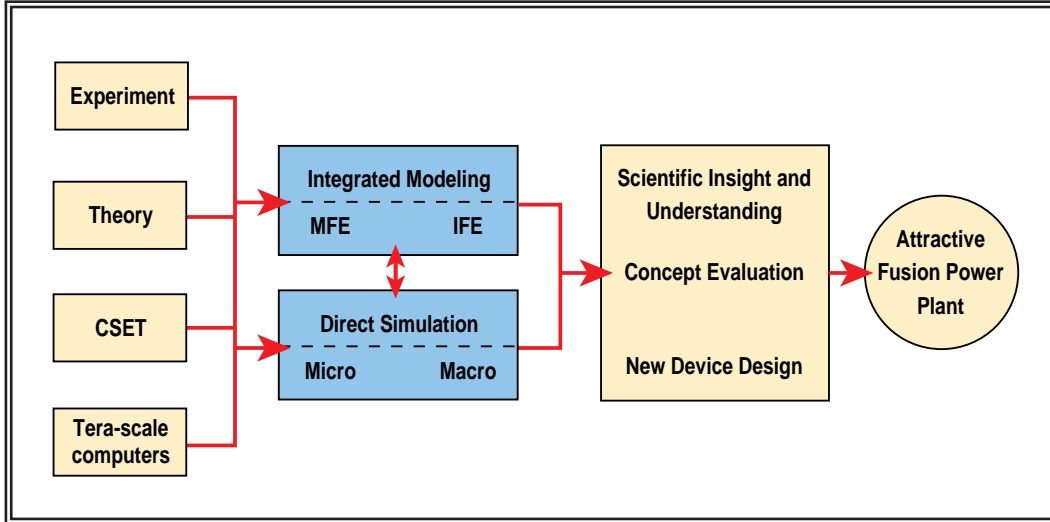


Fig. 4. Fusion SSI elements.

more effectively investigated. It should also be noted that the highly complex plasma systems simulated may be not only more complicated, but its components could be found to interact in ways that produce new, often unexpected effects. The major challenge encountered in the integrated modeling area is to simulate an entire experimental device over macroscopic time scales. All relevant physical processes will be represented at some level by interconnected modules. The new resources made available by the SSI for direct simulation models will enable more effective development of physics-based reduced models that are acceptably realistic and accurate. Hence, the SSI resources would enable integrated modeling at a level of realism and complexity much greater than is presently possible either in MFE or IFE research.

The implementation of the fusion SSI could be cast in three stages which are related to the availability of progressively more powerful high performance computing platforms. It is important to keep in mind here that more powerful computing alone is insufficient to ensure substantive progress without a strong linkage to theory and experiment. Critical scientific issues that can be addressed by 0.5, 5, and 50 teraflop computers and the corresponding increase in memory and data handling capacity are discussed in more detail in the following chapters. As an illustration, a sample roadmap for first principles simulations of plasma microturbulence is depicted in Fig. 5.

Improvements in plasma diagnostic techniques have made it increasingly feasible to demonstrate excellent correlations between experimental results and theoretical models. In particular, the development of diagnostic instruments that can make high resolution measurements of electric and magnetic fields and cross-sectional measurements of turbulent fluctuations have greatly improved the basic understanding of the mechanisms controlling plasma confinement. In order to maximize the effectiveness of simulation/experiment comparisons, it will be necessary to address critical computer science and enabling technology (CSET) issues in the area of data management and visualization. The power of advanced computing to solve critical fusion energy science problems can be fully exploited only if a capable infrastructure is established and effective software tools made available. This can be accomplished if the CSET part of the SSI provides the Fusion Energy Sciences (FES) application with the networking, database, visualization, and other infrastructure tools needed to strengthen the synergistic coupling between simulation and modeling

together with theory and experiment. It is expected that these valuable capabilities will be effectively shared by all the SSI applications. The CSET needs for the FES area is described in Chapter 3 of this document and briefly summarized in Table 1.

An accelerated simulation initiative in fusion energy sciences carries an exciting vision for the future which can serve the vitally

important role of helping to attract and educate the bright young people essential for the future technological health of this field. It would also stimulate scientific alliances with other applications areas in the DOE Office of Science portfolio and increase collaboration with the Accelerated Strategic Computing Initiative (ASCI). Effective modeling of global systems, combustion, and fusion devices all deal with complex, 3D, nonlinear fluid flows and associated kinetic dynamics, albeit in very different parametric regimes. They share the common computational challenge to rapidly develop advanced integrated modeling capabilities capable of treating complex dynamical systems covering many decades in time and space. Some research collaborations already exist between plasma science and ASCI programs, since the physics and mathematical challenges are similar. For example, the computational challenges posed by nonlinear plasma fluid problems share many common features with computational fluid dynamics (CFD) issues faced in ASCI, Global Systems and Combustion Modeling, Materials Sciences, and numerous other areas. As for non-DOE agencies, plasma physics has long-standing programs

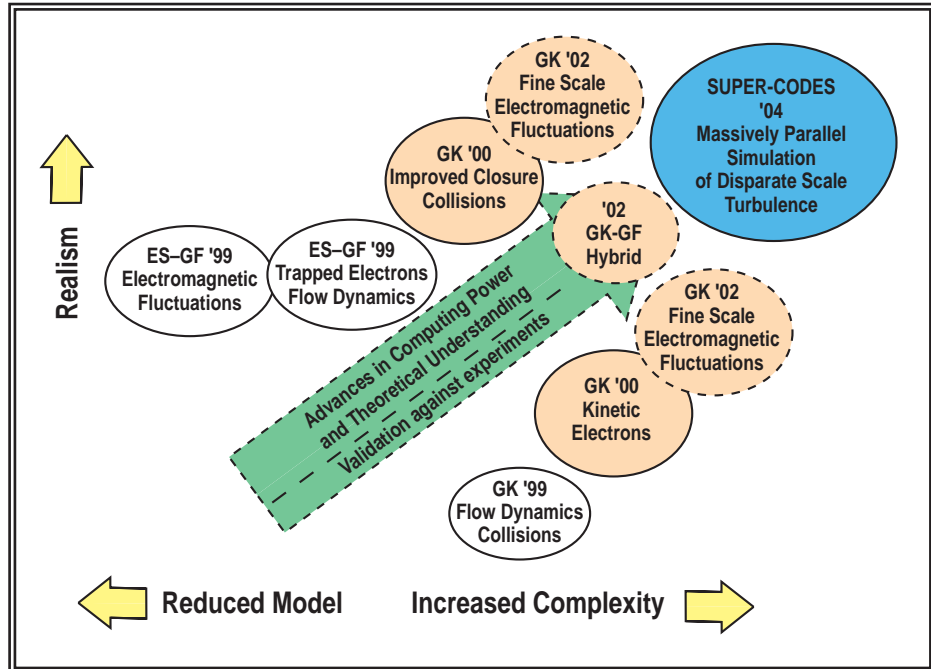


Fig. 5. First principles-based simulations of plasma turbulence.

TABLE 1. CSET NEEDS FOR FUSION SSI

Numerical Algorithms and Libraries	Application Development Environment
<ul style="list-style-type: none"> • Fortran 95 and C++ support • Scaleable parallel differential and matrix solvers • Strongly and partial implicit algorithms • Finite element unstructured meshes • HP computing software libraries 	<ul style="list-style-type: none"> • High-level modular code development and steering techniques • Visualization and image rendering • Code linkage to diagnostics and experimental database • Memory and disk space demands • Distributed computing environment

with NASA and NSF in the space physics and basic science arenas which now could be further strengthened. Additionally, advanced computational modeling capabilities will also stimulate progress in non-fusion plasma applications, including the manufacture of microelectronic components, plasma display panels, plasma thrusters for satellites, and the use of plasmas for waste remediation. Realistic physics-based plasma models will also help accelerate the pace of breakthroughs in plasma science applications to energetic beams, to space physics, to solar physics, and to astrophysics.

In summary, plasma science is prepared to take advantage of the exciting advances in modern computer technology that can transform fusion research in the next decade and accelerate scientific understanding and innovation. New resources in this area will enable productive partnerships involving laboratories, universities, and industries. Sharing with other fields the insights gained in the process of obtaining successful solutions of scientific problems of extraordinary complexity will help to advance all areas of computational science for the benefit of the nation. It will also serve to strengthen the role of the plasma sciences in attracting and training bright young talent needed to ensure the scientific excellence of the field. The advanced computational resources targeted by the present initiative has the potential to truly revolutionize fusion research by improving scientific understanding of experimental data, by stimulating new theoretical ideas, and by providing the most attractive and viable designs for future facilities. This will lead to a more rapid and cost-effective development of the best approaches to fusion energy production; *i.e.*, a “faster” pace toward “smarter” concepts for more practical fusion systems.

2. Critical Science Issues

2.1. Introduction

Plasma is often referred to as the fourth state of matter. A fusion-grade plasma exhibits extremely rich and complex behavior. The unique interplay between the orbits of individual charged particles and the collective effects arising from the long-range nature of the interparticle electromagnetic force lead to a wide range of waves, oscillations, and instabilities which characterize this medium.

The scientific issues related to plasmas generally fall into one of three basic categories: transport, macroscopic stability, and device design. Because charged particles, momentum, and heat move more rapidly along the magnetic field than across it, magnetic fusion research has focused on magnetic traps in which the magnetic field lines wrap back on themselves to cover a set of nested toroidal surfaces called magnetic flux surfaces (because each surface encloses a constant magnetic flux). Transport is concerned with fine-scale turbulence, driven by inhomogeneities in the plasma density and temperature, which can cause particles, momentum, and heat to leak across the flux surfaces from the hot interior to be lost at the plasma edge. Macroscopic stability is concerned with large-scale spontaneous deformations of magnetic flux surfaces. These spontaneous deformations, or macroscopic instabilities, are driven by the large electrical currents flowing in the plasma and by the plasma pressure. Both turbulent transport and macroscopic stability contribute to device design considerations, as does the detailed calculation of particle motions in background electromagnetic fields.

Simulation domains have both minimum and maximum limits on spatial and temporal resolution so that any given plasma simulation can only address a finite range of space and time scales. In the past, this limitation has been dealt with by deriving simplified sets of equations, or “reduced equations,” that are valid for only limited ranges of time and space scales. While the reduced equations have enabled progress in the past, they have fundamental restrictions on their regions of validity. In real plasmas, phenomena occurring on different time and space scales interact and influence one another so that it becomes essential to utilize more general equations, valid on a wider range of space and time scales, and to correspondingly increase our simulation domains.

Two Basic Approaches: Particle-in-Cell and Fluid Models

At the most fundamental level a plasma can be described by kinetic equations for the distribution function within a six-dimensional (plus time) phase space of each particle species. These kinetic equations are coupled to each other through self-consistent electric and magnetic fields. The simulation techniques used in plasma physics fall into one of two broad categories, particle-in-cell models and fluid models. The particle-in-cell methods proceed by integrating a kinetic equation in time by advancing marker-particles along a representative set of characteristics within the phase space. A variety of techniques have been developed over the last 20 years which greatly increase the accuracy and realism of the particle-in-cell simulation technique, while reducing the requirements on the number of “particles” necessary to faithfully represent the physics

The fluid models proceed by advancing velocity moments of the kinetic equation in time. The best known of these are the extended-magnetohydrodynamics (MHD) models which represent the plasma as one or more interacting conducting fluids. This higher-level description frees the model of many fine scale resolution requirements and makes feasible the simulation of large-scale motions and instabilities. Extensive theoretical analysis over the years has led to refinements of the fluid model and improved the closure relations so that many nonfluid effects, such as particle motion and wave-particle resonances, can be represented at some level.

In the following sections we describe representative examples of large-scale simulations that have been performed in the transport, macroscopic stability, and device design areas and discuss future research directions; and the need and readiness for enhanced computational facilities and integrated modeling.

2.2. Microscopic Turbulent Transport Studies

Computer simulation can be used as a test bed to advance the scientific understanding of the turbulence responsible for anomalous transport of particles, momentum, and heat in toroidal magnetic configurations. Representative problems in plasma microscopic turbulence which can be addressed by computer simulation include (1) the role of mesoscales and intermittency, (2) the evolution of transport barriers in the plasma core and at the plasma edge, (3) the role of electron physics and electromagnetic fluctuations, (4) the dynamics of spectral transport, and (5) turbulence in the plasma edge and scrape-off layer. The investigation of these scientific issues will require increasing sophistication in the physics models within our microturbulence codes. In particular, the description of the fields must be improved to include magnetic perturbations (present in some microturbulence codes already) in addition to electrostatic perturbations,

while the description of the electrons must be upgraded to include important kinetic effects like trapping in equilibrium magnetic wells, drift motion, and wave-particle resonances.

The Role of Mesoscales

Toroidal magnetic traps have a minor radius a , which is typically hundreds of times larger than the ion Larmor radius ρ_i . While the dominant turbulence driven by ion temperature gradients (ITG modes) has correlation length typical of ρ_i , there is both experimental [3] and theoretical [4] evidence that ion transport may be dominated by events in which heat (and, presumably, particles and toroidal momentum) propagates quasiballistically across flux surfaces over distances large compared to the radial correlation length. Furthermore, these mesoscale transport events are moderated by self-generated differential rotations of plasma (zonal flows). While this behavior resembles the self-organized criticality paradigm of Bak, Tang, and Wiesenfeld [5], it is significant that the theoretical and computational models of plasma turbulence in which self-similar, scale-invariant behavior has been observed have a firm foundation in the underlying theory. Specifically, are large avalanches dynamically equivalent to radially elongated convective cells (streamers)? Or are they a cascade of localized instabilities driven by the propagating heat pulse? Under what conditions do large events dominate the transport?

The Evolution of Transport Barriers

The best confinement of fusion-relevant plasmas is associated with the spontaneous formation of transport barriers — that is, regions in which the turbulent transport is greatly reduced or vanishes altogether. Transport barriers are associated with regions of high electric field shear [6] and pressure gradients that can approach or exceed ideal MHD stability limits. The understanding and control of these transport barriers is the key problem in reducing the cost of a next-generation DT burning magnetic fusion experiment. In particular, we need to understand the physics behind the experimentally observed power threshold, understand the physics behind the radial propagation of transport barriers, and understand what determines the radial extent of transport barriers. Simulation studies could be used to develop strategies for controlling transport barriers and enabling substantially improved performance in magnetic fusion experiments.

The Role of Electron Physics and Electromagnetic Fluctuations

Plasma microturbulence drives electrical currents parallel to the magnetic field, resulting in magnetic perturbations perpendicular to the equilibrium magnetic field. These magnetic perturbations affect the stability of the underlying microinstabilities, reducing the linear growth rate of the ITG modes while introducing new instabilities, like the kinetic ballooning mode (KBM), associated with deformations of magnetic surfaces; the magnetic perturbations effect the dynamics of zonal flows through the magnetic stress; and they can greatly increase the electron heat transport either by causing transient deformations of the magnetic field or by “breaking” flux surfaces. All of this has a profound effect on microturbulence, as illustrated in Fig. 6. Electromagnetic effects are particularly important in regions of high pressure gradient. Experiments indicate that electron transport can be large within transport barriers where ITG turbulence is absent. This electron transport is likely correlated with the presence of electron temperature gradient (ETG) instabilities. ETG modes have an order of magnitude shorter wavelength than ITC modes, and are inherently electromagnetic, while their proper description requires a kinetic description of the electron species. The resolution of the two wavelengths introduces a significant range of spatial scales,

such that realistic simulations which simultaneously resolve ITG and ETG turbulence will challenge even 50 Tfflop machines.

The Dynamics of Spectral Transport

There is also much to be learned from a Fourier representation of the microturbulence (Fig. 7). Simplified theoretical models have predicted [7], and simulations of drift wave turbulence have observed an inverse cascade. Is this behavior ubiquitous in drift wave turbulence? Is zonal flow excitation a consequence of this phenomenon? There is evidence that the inclusion of a kinetic description of the electrons will at least partially reverse the direction of the cascade [8], perhaps suppressing the self-generated zonal flows. While these issues can be addressed with electrostatic simulations, electromagnetic simulations, including the kinetic electron response, will allow investigation of the modulation interaction between the ITG and the ETG modes. Specialized diagnostics can locate the sources and sinks of spectral energy in k -space, and elucidate the energy transfer physics. Another important issue is that, via nonlinear coupling, long wavelength instabilities can drive electromagnetic transport at small spatial scales by exploiting the collisionless decoupling of the plasma fluid from the confining magnetic field which occurs at spatial scales of the order of the collisionless skin depth (typically about 0.5 mm). This is a challenging question, involving among other issues, the nonlinear transfer of energy over two decades in the spatial scale.




Fig. 6. Contour plot of the electrostatic potential from a sequence of gyrofluid simulations. As the plasma pressure (as measured by $\beta = 2 \mu_0 \rho / B^2$) increases the relatively low frequency ion temperature gradient (ITG) turbulence is replaced by high frequency kinetic ballooning modes.

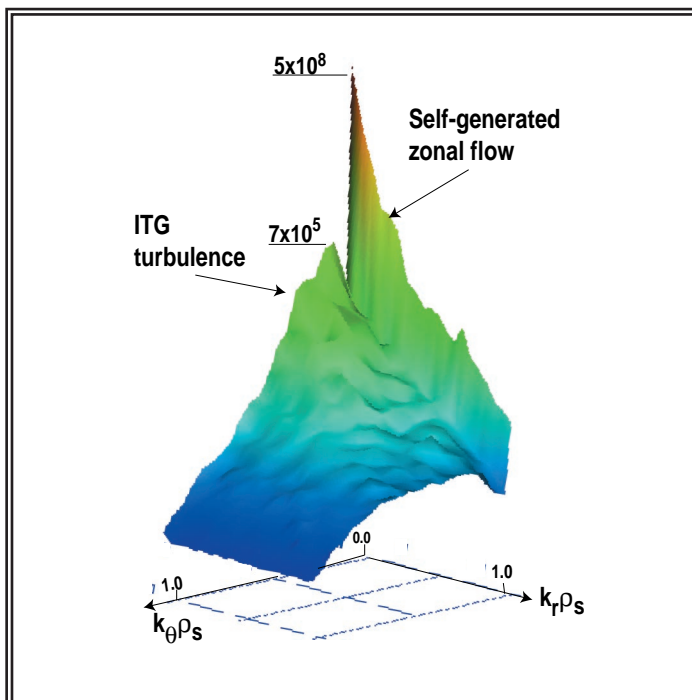


Fig. 7. Spectral density from a gyrokinetic particle-in-cell simulation of ion temperature gradient driven microturbulence.

Flux-Surfaced-Averaged Transport Models

While much work remains to be done to achieve an understanding of microturbulence in magnetically confined plasmas, substantial progress has been made. Flux-surface-averaged models of heat transport have been developed by fitting the ion thermal conductivity to results from nonlinear microturbulence simulations. These models achieve reasonable agreement when compared against actual magnetic confinement experiments. Future improvements to flux-surface-averaged transport models could improve the fidelity of particles and momentum transport, electromagnetic effects in the underlying nonlinear microturbulence simulations, and the resolution of existing discrepancies between the results of various microturbulence simulation models.

Edge Plasma Turbulence

Magnetic confinement plasma terminates at a last closed flux surface. This surface is defined by either the presence of a material boundary, or by a “separatrix” — a magnetic surface separating the closed, nested flux surfaces from open surfaces which strike a material surface. Particles, momentum, and heat are rapidly conducted along open field lines to the materials surface. This results in the formation of a boundary layer about the last closed flux surface where the radial gradients steepen until the cross-field transport can compete with the rapid, along-field-line transport, thus requiring 2D models for equilibrium plasma profiles. Analysis of this boundary layer is complicated by the absence of small parameters, the importance of electromagnetic effects and collisional effects, the presence of neutral gas, and the interaction between the plasma and the material surface. Substantial progress has been made using electromagnetic fluid codes to simulate 3D edge plasma turbulence. Transport barriers controlling core confinement often form near the separatrix in experiments and have been found in these simulations, as have the edge-localized modes which can accompany the formation of such transport barriers.

Future Directions

The proposed R&D effort will be directed at addressing the scientific issues described above. This will require a code development effort which proceeds along three lines: (1) implementation of existing codes on larger computers to enable the simulation of larger problems; (2) enhancements to the physics models in the codes, with particular attention to including electromagnetism, collisional effects and improved electron dynamics; and (3) developing a set of diagnostic and visualization tools that will allow real-time interaction with simulation data, and will assist theorists in testing hypothesis and answering detailed questions regarding the turbulence within these computational models. We anticipate that modern, object-oriented code development methods will facilitate sharing this code development effort among a national team. An important output of this R&D program will be periodic updates to flux-surface-averaged models of turbulent transport to insure that the core transport module in the integrated modeling effort reflects the best current theoretical understanding of plasma turbulence.

2.3. MFE Global Stability Studies

There are many critical scientific problems in fusion science that can be addressed with the macroscopic simulation model known as the extended-MHD model. Most of these share the common features of extreme temporal and spatial stiffness, severe spatial anisotropy, and complex boundary conditions. These characteristics make them among the most challenging problems in

computational physics. Here we describe several examples relevant to large tokamaks, innovative concepts, and basic plasma physics.

High-Pressure Induced Disruptions

One of the most dramatic events to occur in fusion plasmas is known as the major disruption. This is a catastrophic event that can lead to the near instantaneous breakup of the magnetic flux surfaces and loss of confinement. The macroscopic model is capable of predicting the conditions under which the disruption occurs and the consequences of the disruption on the surrounding plasma structure. Figure 2 illustrates the three dimensional structure of select pressure contours during a simulation of a high-plasma-pressure induced disruption in the tokamak TFTR. The disruption mechanism was a complex highly nonlinear process that involved two distinct plasma instabilities interacting to further destabilize one-another nonlinearly. Note that the constant pressure contours resemble fingers in each poloidal plane, but look like ribbons in the toroidal direction. These distortions cause magnetic field lines to become stochastic which in turn causes a rapid loss of thermal energy and a subsequent rapid loss of plasma current to occur.

Nonlinear Evolution of Magnetic Islands

“Magnetic islands” arise from a topological change in the magnetic field mitigated by resistivity or electron inertia. These islands grow on a very slow time scale due to a competition between several effects. They evolve over a time period that, when measured in units of a normalized time for an Alfvén wave to traverse the system, is given approximately by $0.1 S$, where S is the ratio of the magnetic diffusion time to the Alfvén transit time, or the Lundquist number. These islands are the dynamical modes that limit the maximum attainable pressure in long-duration experiments. Figure 8 illustrates the complex magnetic topology structure that can develop in a resistive MHD simulation where plasma instabilities cause the magnetic surfaces to break into islands when the winding numbers on those surfaces are ratios of small integers. This simulation was done with an unrealistic small value of S of 10^4 to ease computational requirements. Computations for realistic values of S , typically of order 10^7 to 10^8 , present an extreme challenge.

Dynamo Activity in a Hot Reversed-Field Pinch (RFP)

The RFP is a magnetic confinement concept which relies on a nonlinear dynamo to sustain the toroidal flux against resistive diffusion. This dynamo arises from the nonlinear interaction of several long wavelength, low frequency dynamical modes. The magnetic fluctuations from these

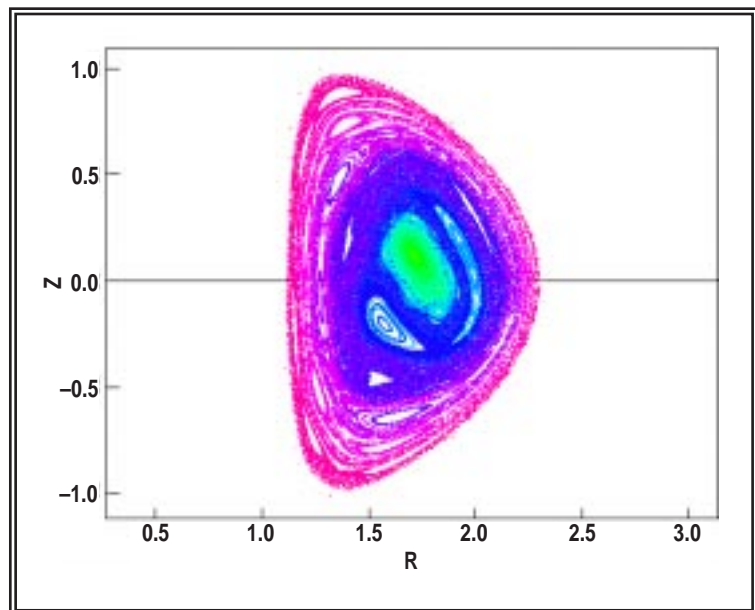


Fig. 8. The complex magnetic topology structure developed as a result of resistive MHD instabilities.

modes determine in large part the confinement properties of this configuration. Critical questions for the future of this concept are how the dynamo activity and the associated magnetic fluctuations scale with S and how this system behaves for times much longer than the resistive decay time of the first wall.

Collisionless Magnetic Reconnection

Changes in magnetic field topology require an electric field parallel to the magnetic field. This electric field can be supported either by finite electric resistivity or by electron inertia. Resistivity is very small in a nearly collisionless plasma, so magnetic reconnection must be a result of electron inertia. Figure 9 illustrates the electron out-of-plane currents in a 2D hybrid (particle ions and fluid electrons) model of magnetic reconnection. The goal is to model this same phenomena in full 3D with both particle ions and electrons. This collision-less reconnection not only occurs in certain fusion experiments, but is also thought to be of fundamental importance in various aspects of solar and space physics, such as solar flares, coronal mass ejection, and magnetic sub-storms in the magnetosphere.

MHD Wave Resonance with Energetic Particles

Plasmas of fusion interest are normally collision-less in the sense that the particle-particle collision mean-free-path is very long compared to all other scale lengths and wave-lengths of interest. This implies that plasma particles that have velocities at or near the wave velocities can interact resonantly with the wave and affect its stability properties. This important effect manifests itself in many ways in fusion plasmas, and is also an important phenomena in space and in astrophysical plasma.

Future Directions

An increase in computing power and improved algorithms will allow 3D nonlinear extended-MHD simulations to include adequate space and time resolution to approach realistic values of the Lundquist number S under fusion conditions, and will allow solving the more complete and more computationally demanding extended-MHD models. For a model of a given mathematical complexity, the computational requirements for 3D extended-MHD scale like S^p where $p = 1-2$, depending on how effective the implicit time-stepping and adaptive mesh refinement are. Therefore, a factor of 100 increase in compute power alone will allow at least an order of




Fig. 9. Electron out-of-plane currents in a 2D hybrid model of magnetic reconnection.

magnitude increase in simulation values, and possibly two orders. Algorithmic advances and better single-processor optimization may lead to another order of magnitude. Thus, we can expect that the fluid models, now limited by computational resources to compute at S values in the range of 10^5 to 10^6 , to be able to approach simulating plasmas with fusion relevant values of 10^7 to 10^8 . The increase in computing resources will have an equally dramatic effect on enabling the codes utilizing the physical models that involve particle closures to resolve velocity space. As we are able to exercise these computationally demanding models with better resolution and for longer times, we will develop better understandings of the physical differences between fluid and particle closures, which we can expect to lead to improved and more efficient models.

2.4. IFE Source-to-Target Ion-Beam Simulations

The application of inertial confinement principles to electric power production requires the successful development of an efficient driver. Following the recommendation of several advisory panels [9,10], the driver-development effort within the Office of Fusion Energy Sciences has focused on heavy ion accelerators as drivers for Inertial Fusion Energy (IFE) systems. Heavy ion drivers offer the prospect of high efficiency, long lifetime, flexible target illumination geometry, high repetition rate, and a robust magnetic final focusing lens. An end-to-end heavy-ion driver simulation requires the linkage of three classes of models: (a) particle simulations of the ion beam dynamics from their source to the exit of the accelerating and pulse-compressing system, (b) large-scale electromagnetic (or magnetoinductive) simulations of ion beam propagation through the fusion chamber, and (c) detailed microscopic simulations which can capture electron effects and collective modes in the driver and fusion chamber, to ensure that these do not disrupt the focusability of the beams. This strategy for enhancing the computational support for heavy ion fusion development is illustrated in Fig. 10.

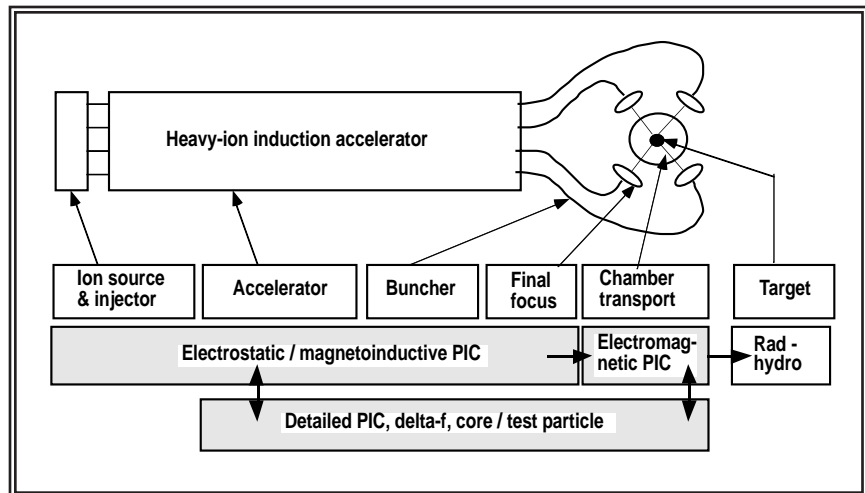


Fig. 10. Schematic of a heavy-ion based inertial fusion energy (IFE) system. The models to be developed and integrated are indicated by the shaded boxes.

If a heavy-ion driver is to be cost-effective, it must accelerate as large an ion current density as possible. As a result, the intense beams in a heavy-ion driver are *non-neutral plasmas* whose dynamics are dominated by space-charge effects (rather than by thermal pressure as in most accelerators) [11,12]. It is essential that “emittance growth” (beam phase-space dilution) and “halo generation” (creation of an outlying beam particle population) be minimized. These effects are kinetic in nature, and can be addressed with particle simulation techniques analogous to those used to study plasma microturbulence. Sophisticated 3-D particle-in-cell simulations of space-charge-limited ion beams have already been carried out for sections of a heavy-ion driver

(see Fig. 3), and 2-D calculations (which assume steady, paraxial flow) have simulated complete near-term accelerators. However, effective accelerator design requires integrated 3-D end-to-end modeling. Beam perturbations in one accelerator section affect the particle distribution function in the next section, ultimately limiting the achievable beam focus at the target. This calculation is made more difficult by the fact that low probability events contribute, *e.g.*, to halo formation. Furthermore, stray electron effects in the driver need to be studied, in simulations which employ timesteps considerably smaller than those set by the ions alone. Present simulations cannot model the entire length of a driver-scale accelerator in full 3-D, although the algorithmic and hardware enhancements necessary are fairly well understood.

The results from the accelerator must be carried into the final compression line, final-focus optical system, and target chamber in a self-consistent way. The beam has a time-varying energy and transverse particle distribution, while the target design requires a particular pulse shape (current vs. time) at the final optic. Thus the optical aberrations will be time-varying, and must be controlled. Time-varying current effects in the chamber affect the focusing, and must be modeled consistently with the partial beam neutralization in the target chamber. These problems have been examined to a limited extent, but new computer capabilities and now algorithms to increase the spatial and temporal scales in simulations are needed to include multi-beam effects, inductive effects at the chamber entrance, and improved atomic physics models.

2.5. Readiness for Terascale Computing

MFE Microscopic Simulations

The fusion community has been very successful in developing codes for modeling plasma turbulence for which computer run-time and problem size scale well with the number of processors on massively parallel machines [2] (see Fig. 11).

This experience provides the basis for our confidence that we will be able to obtain significant results during the first year of operation of the SSI computers. To illustrate this, we consider a full-device simulation of a next-step tokamak capable of achieving ignition or high fusion gain. Such a device would have a minor radius of the order of $10^3 \times \rho_i$. Using a magnetic field aligned coordinate system, a full-device simulation of ITG turbulence, including wavelengths down to the ρ_i scale, would require a grid with $N_{\text{Radial}} \approx 1000$, $N_{\text{Poloidal}} \approx 3000$, and $N_{\parallel} \approx 64$ for a total of about 2×10^8 grid points. Our experience indicates that eight particles/grid cell (for a total of 1.6×10^9 particles) is adequate to provide a long interval of fully developed turbulence without excessive particle noise. About 50,000 time steps will be required to simulate 100 turbulent decorrelation times.




Fig. 11. 3D gyrokinetic particle-in-cell codes have demonstrated excellent scaling as the number of processors is increased to 512 (the maximum available).

Scaling from our experience on the T3E (a 0.5 TFlop machine) we find that a 5 TFlop computer can be expected to achieve a 4×10^{-9} sec/particle/timestep. Hence, a full-device simulation of ITG turbulence an ignition experiment will take about 90 hours on the SSI computer. Presently, our codes require about 0.3 kbyte of memory/particle. However, the larger grid required for these simulations will may require domain-decomposition in two or three dimensions (presently we only decompose in 1D), leading to somewhat larger memory requirements/particle. Taking a conservative estimate of 1 kbyte/particle, we conclude that we this simulation would require about 1.6 Tbyte of RAM distributed among the processors. The disk storage requirements, which include data for diagnostics and restart dumps, would be tens of Tbytes.

MFE Macroscopic Simulations

Several 3D Nonlinear macroscopic simulation codes have been developed within the fusion program that have already demonstrated good scaling on parallel computers for problems of interest to this project. Examples of these are briefly described.

One of the major macroscopic fusion plasma simulation codes [13], solves the nonlinear, 3D two-fluid MHD evolution equations in toroidal geometry using a strongly implicit time advancement algorithm that allows the dominant linear resistive MHD response to be solved each time step in a strongly coupled form. The electron inertia and Hall terms are also treated in an implicit manner. This strong implicit treatment decouples the time step from the fastest time scales and allows the time step to be much larger than the transit time of an Alfvén wave, and in principle to scale like $S^{1/2}$. Spatial approximation in this code is by means of combining grid blocks of either structured (rectangular) or unstructured (triangular) finite elements in the poloidal plane, and fast Fourier transforms (FFTs) in the toroidal direction. From its design inception, the computationally intensive physics kernel has been developed so as to enable a parallel implementation of the code. The grid blocks allow efficient decomposition of the poloidal plane. In addition, the form of the implicit algorithm leads to separate matrix equations for the different Fourier components; coupling occurs only through explicit nonlinear terms. This permits advancing groups of Fourier components on separate layers of processors, where each layer may have a number of processors addressing different grid blocks. Thus, all three spatial dimensions of a simulation domain may be assigned to different processors in a distributed fashion. Standard message-passing library routines (MPI) are used to communicate data between processors without the need for hardware support of shared memory. The code allows for two types of decomposition (grid blocks and Fourier groupings) to be performed in tandem, and the resulting parallel speed-up is essentially the product of the speed-up due to each decomposition acting separately. This results in nearly ideal performance scaling on massively parallel computers for nonlinear implicit problems.

Another of the major 3D extended MHD codes [14] utilizes a partially implicit approach for its time advancement. Each time step has both fluid and particle time advancement. For the fluid part, 1-D (toroidal direction) parallelization has been done on a 64 processor Origin2000 using OpenMP and good scalability has been obtained. The particle advancement is completely parallel. This code uses a stream function/potential representation for the magnetic vector potential and velocity that has been designed to minimize spectral pollution. The basic solution algorithm is quasi-implicit in that only certain terms in the fluid part that are the most timestep limiting are solved for implicitly, with explicit differencing being used for the remaining terms. Each timestep, several uncoupled 2-D scalar elliptic equations are be solved within each poloidal

plane using an incomplete Cholesky decomposition with the conjugate gradient method, with the different poloidal planes being computed concurrently. The parallelization in this code is now being extended to multiple dimensions as required for massively parallel distributed memory computers. Work is continuing in optimization of the iterative matrix solver, and the scalability of different parallel solvers is being evaluated. The further development and optimization of hybrid particle-in-cell/fluid closure schemes appropriate for long-mean-free-path plasmas is also an active area of research, and idealized boundary conditions are being replaced by more realistic models to better represent physical configurations of interest.

IFE Systems

A key 3-D particle-in-cell code [15] which combines plasma particle-in-cell techniques with detailed descriptions of the elements of an accelerator has been developed, and is in use at LLNL, LBNL, NRL, and The University of Maryland. It uses a 3-D Cartesian mesh which naturally represents a linear self-field variation with transverse position when Poisson’s equation is solved with a uniform charge density. This is important because a field (applied plus space charge) which varies linearly does not lead to emittance growth, but any representational errors would lead to spurious growth. Furthermore, the simplicity of this mesh keeps the particle advance efficient. The field-solver and particle advance have been designed to handle a “warped” coordinate system, so that accelerators which incorporate bends can be modeled naturally. Each element of the accelerator is described in a coordinate system natural to itself, *e.g.*, with the local *z*-axis passing down its centerline. User-programmable code steering is incorporated via the public-domain Python interpreter.

This code runs in a parallel mode on existing supercomputers such as NERSC’s Cray T3E-900. Spatial domain decomposition techniques are used, and the speed of the particle advance scales extremely well with the number of processors. The solution of Poisson’s equation requires global communication across processors, but good parallel scaling of FFT and SOR methods has been obtained using up to 256 processors, the most tested to date. The problems to be tackled under the SSI will require a more detailed applied field description, and scalability to thousands of processors is expected.

For studies of beam propagation through the fusion chamber environment and onto the target, electromagnetic particle codes simulating the beam and the background plasma ions and electrons as discrete particles in converging 2-D geometry, and recently in full 3-D geometry, are strong starting points for the needed capability. These codes have been exercised to explore some scenarios, but have not yet treated the important and far more challenging effects associated with pulse-shaped beams, multi-beam interactions, and dense and inhomogeneous background plasmas.

A 3-D multispecies nonlinear-perturbative [16] particle code is under development for detailed studies of intense beam propagation and stability, and beam-plasma interaction processes. This code should adapt readily to terascale computers.

2.6. Readiness for Integrated Modeling — MFE Devices

Integrated modeling is essential to the mission of controlled fusion research. While specialized computations can explore turbulence and large scale instabilities in isolation, only integrated modeling can bring all the strongly interacting physics together to make coherent, self-consistent predictions about fusion devices.

The integrated simulation of a magnetic fusion device involves the simultaneous modeling of the core plasma, the edge plasma, and the plasma-wall interactions. Integrated simulations have a major impact on the design of new magnetic fusion devices and ultimately on the design of a fusion power plant. In each region of the plasma, there is anomalous transport driven by turbulence, there are abrupt rearrangements of the plasma caused by large-scale instabilities, and there are interactions with neutral atoms and electromagnetic waves. Many of these processes must be computed on short time and space scales, while the results of integrated modeling are needed for the whole device on long time scales. The mix of complexity and widely differing scales in integrated modeling results in a unique computational challenge, one that is on the scale of modeling climate and combustion. Building on the experience of the National Transport Code Collaboration (NTCC) [17,18], this initiative will develop a framework for complex model integration and to use this framework to couple components in both core and edge regions of a fusion device.

An integrated modeling framework will be developed which implements parallel and simultaneous calculations of logically distinct physics components, using objects of variable granularity. The framework would gather timing information so that it could control the granularity of the separate calculations, combining fast calculations onto single processors, and requiring slow calculations to granularize further, as it steps through the calculation. This process is called “heterogeneous parallelization,” where data parallelism (*e.g.*, domain decomposition) and functional parallelism (process decomposition) are used together in an optimal, problem-dependent manner [19]. This machinery must be eventually capable of handling a wide range of object complexity concurrently, ranging from multi-dimensional components to many simultaneous relatively small calculations. In parallel, this project will develop the networking capability for job monitoring and even steering. This will require periodic launching of processes to extract data for visualization and presentation of that data on client workstations [20].

There are some modules in integrated modeling codes that would clearly benefit from massively parallel computations. For example, Monte Carlo computations are used extensively to follow neutrals and complex atomic physics in the plasma [21] and to compute the plasma interaction with the material wall which is needed to determine erosion rates, recycling, and plasma contamination [22]. These computations as well as ray tracing for radio frequency and radiation transport codes are relatively easy to parallelize. As the integrated modeling framework evolves, it will include more modules that are increasingly computationally challenging. For example, more sophisticated edge transport and turbulence using fluid and kinetic models will be incorporated [23, 24]. This will enable more detailed comparisons with experimental data and will yield more reliable integrated modeling predictions.

3. Computer Science and Enabling Technology (CSET) Needs and Development

The SSI can benefit extensively from the experience of the High Performance Computing and Communication (HPCC) Grand Challenges and the ASCI. These initiatives have recognized from the outset that a truly productive Tera-scale computing resource requires a complementary software resource that scales as well as the hardware and which provides a rich scalable code

development environment. Significant funding and manpower have been invested in CSET development to meet these objectives and it offers a valuable resource that SSI can tap into.

The CSET needs for fusion SSI applications fall into two general categories: (1) *Numerical Algorithms and Software Libraries* — in SSI, the researchers will be tackling scientific problems with such complexity that, in addition to making sure presently working algorithms are scalable to new computing platforms, innovative numerical algorithms will have to be invented to make progress. (2) *Application Development Environment* — a critical challenge for SSI will be that applications development will involve creation of large, complex codes by multiple institutions. Present development environments for scientific software are antiquated and not well suited to the development of re-usable, maintainable, and extensible codes, particularly by distributed development teams. Sophisticated high-level software development environments are needed to support effective use of terascale machines. Furthermore, to ensure the fidelity of the science, it is essential for the applications to be rigorously validated against experiments. Connectivity to experimental data and state-of-the-art tools for data visualization, mining and manipulation will be essential. The applications communities do not generally have the resources to provide their own CSET needs; nor is it the most efficient way, since many of the CSET tools have cross-cutting benefits for multiple applications. An integrated effort led by the CSET community, working in partnerships with the applications areas, HPCC and ASCI. will be required.

Fusion SSI CSET needs have been integrated in the discussions of critical scientific issues in Chapter 2. Below, key needs are highlighted.

Numerical Algorithms and Libraries

- Most of the existing fusion codes are Fortran-based, and are already being converted to Fortran 95. Good support on the SSI Tera-scale computers for both Fortran 95 and C++ will be essential.
- For particle simulations, research into elliptic equation solvers that scale well to thousands of processors would have a significant payoff.
- For macro-scale simulations, the challenge in simulating many decades of time scale and a hierarchy of physics levels suggests further considerations of both strongly implicit and partially implicit algorithms.
- Finite element unstructured meshes work well in resolving multi-scale spatial structures, encountered in many problems of increasing realism, but requires further improvement in multi-dimensional parallelization.
- Further work is required in optimization of the iterative matrix solver, including sparse matrix solver for fluid codes; and the scalability of different parallel solvers should be evaluated.

Application Development Environment

- Considerable development is needed to reap the full benefit from user-programmable code-steering techniques and the exploitation of high-level modular code development tool kits.
- Means of dealing with the “data glut” in the interactive exploratory visualization of Tera-scale simulations, including image rendering, must be developed.

- Convenient means for linking code results into the diagnostic system and shot database from physics experiments, and for folding experimental results into simulation diagnostic output for comparison purposes, would have a high payoff.
- Memory requirements, though less demanding than CPU requirements, still requires attention. Commensurate disk space, allowing for multiple full data dumps during a long run and multiple runs archived, is also needed.
- System management tools to enhance the robustness of the code development environment including Distribute Computing Environments will be highly beneficial.

References

- [1] President's Committee of Advisors on Science and Technology, "Report to the President on Federal Energy Research and Development for the Challenges of the Twenty-First Century," November 1997.
- [2] Z. Lin *et al.*, *Science* **281**, 1835 (1998); A.M. Dimits *et al.*, *Phys. Rev. Lett.* **77**, 71 (1996); J. Kepner *et al.*, *SIAM News* **30**, 1 (1997); R.E. Waltz *et al.*, *Phys. Plasmas* **4**, 2482 (1997).
- [3] R.A.Moyer *et al.*, *Bull. Am. Phys. Soc.* **43**, 1664 (1998); P.A. Politzer *et al.*, *ibid.*; B.A. Carreras, *ibid.*
- [4] B.A. Carreras *et al.*, *Phys. Plasmas* **3**, 2903 (1996); X. Garbet *et al.*, *Phys. Plasmas* **5**, 2836 (1998).
- [5] P. Bak *et al.*, *Phys. Rev. Lett.* **59**, 381 (1987).
- [6] K.H. Burrell, *Phys. Plasmas* **4**, 1499 (1997).
- [7] A. Hasegawa and K. Mima, *Phys. Rev. Lett.* **39**, 205 (1977); A. Hasegawa and K. Mima, *Phys. Fluids* **21**, 87 (1978).
- [8] L. Chen *et al.*, *Phys. Rev. Lett.* **39**, 754 (1977).
- [9] Fusion Policy Advisory Committee Final Report, Sept. 1990, DOE/S-0081.
- [10] Fusion Energy Sciences Advisory Committee Report, July 1996, DOE/ER-0690.
- [11] R.C. Davidson, *Phys. Rev. Lett.* **81**, 991 (1998).
- [12] A. Friedman, Proc. IUPAP / Int. Committee for Future Accelerators Workshop on Space Charge Dominated Beams and Applications of High Brightness Beams, Bloomington IN, AIP Conf. Proc. **377**, 401 (1996), S. Y. Lee, Ed.
- [13] A.H. Glasser *et al.*, in *Plasma Phys. Contr. Fusion*, 1999.
- [14] "Plasma Simulation Studies using Multilevel Physics Models," W. Park, W.V. Belova, G.Y.Fu, X. Tang, H.R. Strauss, L.E. Sugiyama, to appear in *Phys. Plasmas* (1999).
- [15] A. Friedman *et al.*, *Phys. Fluids B* **4**, 2203 (1992).
- [16] P.H. Stoltz *et al.*, *Phys. Plasmas* **6**, 298 (1999).
- [17] J.R. Cary *et al.*, "Client-Server for Remote Collaboration and Remote Invocation of the Analysis Applications and Application to the National Transport Code," Proc. Workshop, Remote Participation in Fusion Experiments (Prague, Czech Republic, 1998) V. Piffel and J. Pichal, Eds., ISBN 80-01-01829-6.
- [18] A.H. Kritz *et al.*, "Remote Participation in the National Transport Collaboration," Proc. Workshop on Remote Participation in Fusion Experiments (Prague, Czech Republic, 1998) V. Piffel and J. Pichal, Eds., ISBN 80-01-01829-6.
- [19] http://www.tc.cornell.edu/Edu/Talks/IntroToPP/less.html#section_4 (Functional and data parallelism).
- [20] R. Orfali and D. Harkay, *Client/Server Programming with Java and CORBA* (John Wiley & Sons, 1997).
- [21] D. Stotler and C. Karney, *Contrib. Plasma Phys.* **34**, 392 (1994). <http://w3.pppl.gov/degas2/>.
- [22] J.N. Brooks, in *Atomic and Plasma-Material Interaction Processes in Controlled Thermonuclear Fusion*, R.K. Janev and H.W. Darwin, Eds. (1993) Elsevier Science Publishers B.V., p. 403. J.N. Brooks, D.G. Whyte, "Modeling and Analysis of DIII-D/ DiMES Sputtered Impurity Transport Experiments," to be published in *Nuclear Fusion*.
- [23] X.Q. Xu *et al.*, *Contr. Plasma Phys.* **36**, (1998).
- [24] T. Rognlien *et al.*, *Contr. Plasma Phys.* **34**, 362.