

## TOPICAL REVIEW

# Advances and challenges in computational plasma science

**W M Tang<sup>1</sup> and V S Chan<sup>2</sup>**<sup>1</sup> Princeton Plasma Physics Laboratory, Princeton University, PO Box 451, Princeton, NJ 08543, USA<sup>2</sup> General Atomics, PO Box 85608, San Diego, CA 92138, USA

Received 1 June 2004, in final form 21 October 2004

Published 11 January 2005

Online at [stacks.iop.org/PPCF/47/R1](http://stacks.iop.org/PPCF/47/R1)**Abstract**

Scientific simulation, which provides a natural bridge between theory and experiment, is an essential tool for understanding complex plasma behaviour. Recent advances in simulations of magnetically confined plasmas are reviewed in this paper, with illustrative examples, chosen from associated research areas such as microturbulence, magnetohydrodynamics and other topics. Progress has been stimulated, in particular, by the exponential growth of computer speed along with significant improvements in computer technology. The advances in both particle and fluid simulations of fine-scale turbulence and large-scale dynamics have produced increasingly good agreement between experimental observations and computational modelling. This was enabled by two key factors: (a) innovative advances in analytic and computational methods for developing reduced descriptions of physics phenomena spanning widely disparate temporal and spatial scales and (b) access to powerful new computational resources. Excellent progress has been made in developing codes for which computer run-time and problem-size scale well with the number of processors on massively parallel processors (MPPs). Examples include the effective usage of the full power of multi-teraflop (multi-trillion floating point computations per second) MPPs to produce three-dimensional, general geometry, nonlinear particle simulations that have accelerated advances in understanding the nature of turbulence self-regulation by zonal flows. These calculations, which typically utilized billions of particles for thousands of time-steps, would not have been possible without access to powerful present generation MPP computers and the associated diagnostic and visualization capabilities. In looking towards the future, the current results from advanced simulations provide great encouragement for being able to include increasingly realistic dynamics to enable deeper physics insights into plasmas in both natural and laboratory environments. This should produce the scientific excitement which will help to (a) stimulate enhanced cross-cutting collaborations with


other fields and (b) attract the bright young talent needed for the future health of the field of plasma science.

## 1. Introduction

Often referred to as ‘the fourth state of matter’, plasmas comprise over 99% of the visible universe and are rich in complex, collective phenomena. The quest for harnessing fusion energy is a major component of research in the field of plasma physics. Fusion is the power source of the Sun and other stars, which occurs when forms of the lightest atom, hydrogen, combine to form helium in a very hot ( $10^8$ °C) ionized gas or ‘plasma’. The development of fusion as 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. As such, progress toward this goal requires the acquisition of the basic scientific understanding to enable the innovations that are still needed for making fusion energy practically realizable. In this, as well as other areas facing major scientific challenges, it is well recognized that research in plasma science requires the accelerated development of computational tools and techniques that aid the acquisition of the scientific understanding needed to develop predictive models which can prove superior to empirical scaling. This is made possible by the rapid advances in high-performance computing technology, which will allow simulations of increasingly complex phenomena with greater physics fidelity. Accordingly, advanced computational codes, properly benchmarked with theory and experiment, are now generally recognized to constitute a powerful new tool for scientific discovery. In this paper, recent progress and future directions for advanced simulations in magnetically confined plasmas are reviewed. Illustrative examples from magnetic confinement investigations, which include magnetohydrodynamics (MHD), microturbulence, magnetic reconnection and other topics, have largely been drawn from the work of US research groups. It should be noted that important progress has been made worldwide and that this review does not constitute a comprehensive survey with associated references of all of the significant advances in computational plasma physics. The purpose here is to highlight the progress and remaining challenges in those applications which illustrate that plasma science is effectively utilizing and contributing to the exciting progress in information technology and scientific computing [1]. Overall, this area of research in plasma science has helped create a more stimulating environment for carrying out the type of research with the greatest promise for accelerating scientific understanding, innovation and new discoveries.

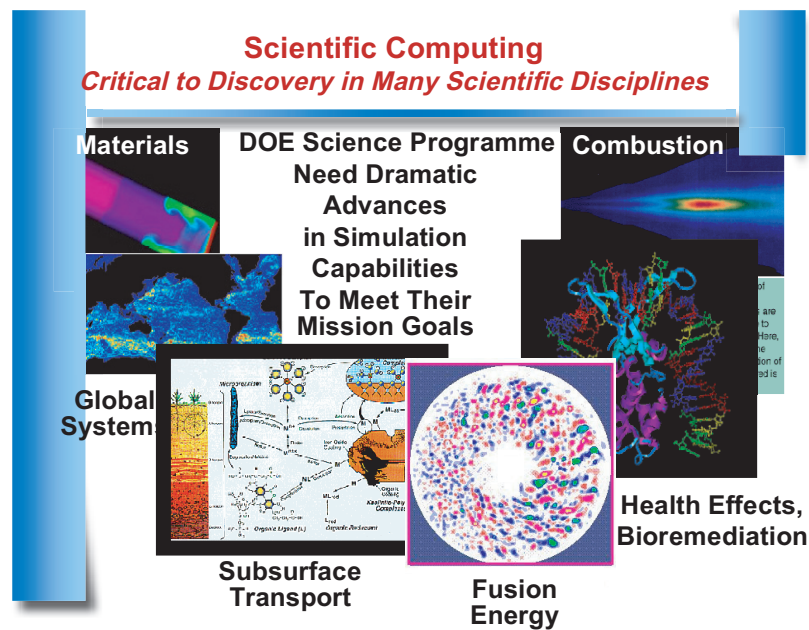
## 2. Challenges for plasma theory and simulations

In a magnetically confined plasma, the interplay between the complex trajectories of individual charged particles and the collective effects arising from the long-range nature of electromagnetic forces leads to a wide range of waves, oscillations and instabilities characterizing the medium. As a result, there is an enormous range of temporal and spatial scales involved in plasmas of interest. As illustrated in figure 1, the relevant physics can span over ten decades in time and space. Associated processes include the turbulence-driven (‘anomalous’) transport of energy and particles across a confining magnetic field, the abrupt rearrangements (disruptions) of the plasma caused by large-scale instabilities, and the interactions involving the plasma particles with electromagnetic waves and also with neutral atoms. Many of these phenomena involve short length and time scales (nanoseconds and microns) while others occur on long time scales (seconds and minutes) and length scales



**Figure 1.** Huge ranges in spatial and temporal scales present major challenges to plasma theory and simulation.

in the order of the device size (metres). Although the fundamental laws that determine the behaviour of 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. Effective prediction of the properties of energy-producing fusion



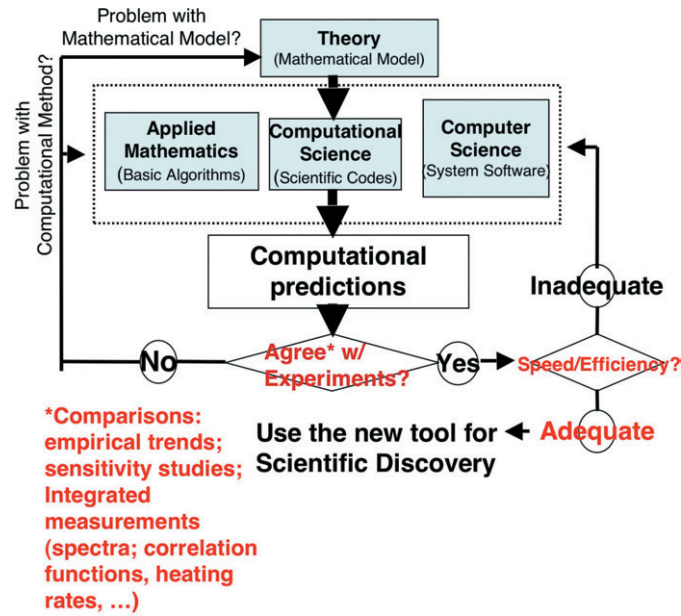
**Figure 2.** Scientific Computing is critical to discovery in many scientific disciplines such as fusion energy sciences.

plasma systems depends on the successful integration of many complex phenomena spanning vast ranges. This is a formidable challenge that can only be met with advanced scientific computation properly cross-validated against experiment and analytic theory.

In magnetic confinement fusion experiments, the plasma interacts directly with the ‘confining’ electromagnetic fields, which can come from an external source and/or from currents produced within the plasma. This can lead to unstable behaviour, where the plasma rapidly rearranges itself and relaxes to a lower energy state. The resultant thermodynamically favoured state is incompatible with the conditions needed for fusion systems, which require that more power output be generated than it takes to keep the hot plasma well confined. However, the hot plasma state is naturally subject to both large- and small-scale disturbances (‘instabilities’) which provide the mechanisms for lowering its energy state. It is therefore necessary to first gain an understanding of these complex, collective phenomena, and then to devise the means to control them. The larger-scale (‘macro’) instabilities can produce rapid topological changes in the confining magnetic field resulting in a catastrophic loss of fusion power density. Even if these instabilities can be controlled and/or prevented, there can remain smaller-scale (‘micro’) instabilities which prevent efficient hot plasma confinement by causing the turbulent transport of energy and particles. In order to make progress on these issues, researchers in this field have effectively utilized advanced computation and modelling over the years, to deal with the complexity of the kinetic, electromagnetic and atomic physics equations describing the behaviour of fusion plasmas. Recognition of the vital role of advanced simulation of plasmas has been underscored by its inclusion as a prominent element of the US Department of Energy’s ‘Scientific Discovery through Advanced Computing’ (SciDAC) programme [2]. As shown in figure 2, the SciDAC programme has emphasized that advanced computing is critical to discovery across many science disciplines [3]. The Fusion SciDAC projects include many of the key topical areas discussed in this review.

The scientific challenges related to magnetically confined plasmas can be categorized into four areas: macroscopic stability, wave–particle interactions, microturbulence and transport and plasma–material interactions. In addition, the integrated modelling of the physical processes from all of these areas is needed to effectively (a) harvest the physics knowledge from existing experiments and (b) design future devices. 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). Macroscopic stability is concerned with large-scale spontaneous deformations of magnetic flux surfaces. These major displacements or macro-instabilities are driven by the large electrical currents flowing in the plasma and by the plasma pressure. Wave–particle interactions deal with how particles and plasma waves interact. Detailed calculations of particle motions in background electromagnetic fields are needed, in order to assess the application of waves to heat the plasma as well as address the dynamics of energetic particles resulting from intense auxiliary heating and/or alpha particles from the fusion reactions. Microturbulence and the associated transport come from 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. Plasma–material interactions determine how high-temperature plasmas and material surfaces can co-exist. Progress in the scientific understanding in all of these areas contributes in an integrated sense to the interpretation and future planning of fusion systems. This demands significant advances in physics-based modelling capabilities—a formidable challenge which highlights the fact that advanced scientific codes are a realistic measure of the state of understanding of all natural and engineered systems.

As illustrated schematically in figure 3, the path for developing modern high-performance computational codes as validated tools for scientific discovery involves a multi-disciplinary collaborative process. This begins with basic theoretical research laying the foundations for the mathematical formulation of the physical phenomena of interest observed in experiments. Computational scientists produce the codes that solve these equations, using the best possible algorithms that efficiently utilize modern high-performance computers. They do so in partnership with applied mathematicians, who provide the basic mathematical algorithms, and the computer scientists, who provide the requisite computer systems software. The computational scientists must then engage the research applications scientists in the critical scientific code validation phase, where the newly computed results are compared against experimental/observational data. This is a major challenge involving a hierarchy of benchmarking criteria which begin with cross-checks against analytic theory, empirical trends and suggestive ‘pictorial’ levels of agreement. It then graduates to sensitivity studies, where agreement is sought when key physical parameters are simultaneously varied in the simulation and experiment/observation. At the next level, richer physics validation is dependent on the availability of advanced experimental diagnostics which can produce integrated measurements of key physical quantities such as spectra, correlation functions, heating rates and other variables of interest. If the simulation/experimental data comparisons are unsatisfactory at any of these validation levels, the work flow moves back to: (a) the theorists (in consultation with experimentalists) if the problem appears to be with the mathematical model and (b) computational scientists (in consultation with applied mathematicians and computer scientists) if the problem appears to be with the computational method. Even when the theory/experiment comparisons prove satisfactory, code performance criteria for speed and efficiency could dictate another round in the computational science box. If all criteria are met, then the new ‘tool for scientific discovery’ can be effectively utilized for interpreting



**Figure 3.** Development path for high-performance codes as validated tools for scientific discovery.

experimental data, designing new experiments and even predicting new phenomena of interest. This cycle of development will of course be repeated as new discoveries with associated modelling challenges are encountered. In addition, it should be kept in mind that the continuous development of a robust computational infrastructure (including hardware, software and networking) is needed to enable capabilities which minimize time-to-solution for the most challenging scientific problems.

It is clear that since any given plasma simulation can only address a finite range of space and time scales, the associated domains have both minimum and maximum limits on spatial and temporal resolution. Accordingly, simulation models are commonly developed from simplified sets of equations, or ‘reduced equations’, which are valid for only limited ranges of time and space scales. Examples include ‘gyro-kinetic equations’ [4] for dealing with turbulent transport problems and the ‘MHD equations’ [5] for addressing the large-scale stability issues. While the reduced equations have enabled progress in the past, fundamental restrictions on their regions of validity have motivated the drive for improvements. In actual laboratory or natural plasmas, phenomena occurring on different time and space scales interact and influence one another. Simulations with greater physics fidelity thus demand increased simulation domains, which can only result from the derivation and application of more general equations that are valid on a wider range of space and time scales.

The most fundamental theoretical description of a plasma comes from kinetic equations for the distribution function within a six-dimensional phase space of each particle species (plus time). They are coupled to each other through self-consistent electric and magnetic fields. Velocity moments of these kinetic equations produce a hierarchy of fluid equations amenable to modelling. In general, the simulation techniques used in plasma physics fall into two broad categories: kinetic models and fluid models. The most mature kinetic approach is the particle-in-cell method, pioneered by John Dawson and others [6]. This method involves integrating a (possibly reduced) kinetic equation in time by advancing marker particles along a representative set of characteristics within the (possibly reduced) phase space. It basically

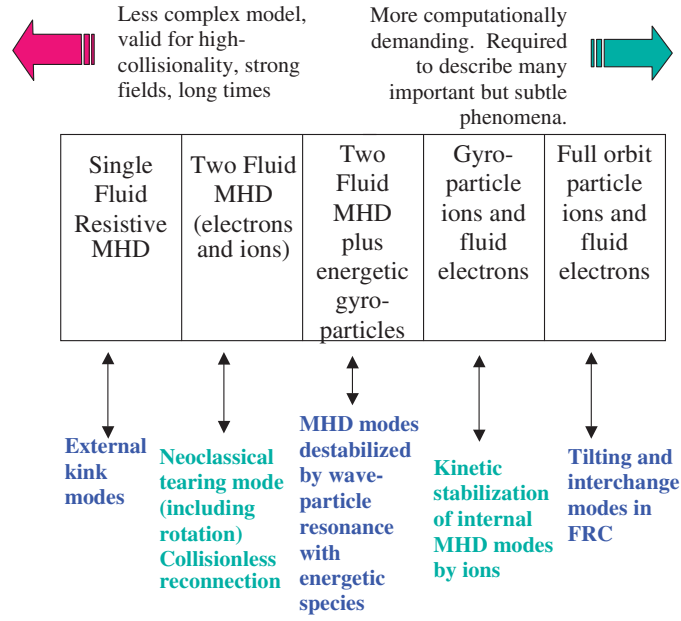
involves a Lagrangian formulation in which the dynamics of an ensemble of gyro-averaged particles are tracked. Simulation techniques such as ‘finite sized particles’ [7] (to reduce the ‘noise’ coming from discrete marker particles), ‘gyro-kinetics’ [8] (a reduction of the full kinetic equation to a five-dimensional phase space which removes high-frequency motion that is not important to turbulent transport) and ‘delta-f’ [9] (a prescription for further reducing the discrete particle noise by separating the perturbed from the equilibrium part of the distribution function before integrating the gyro-kinetic equation along the appropriate characteristics) have been developed over the last 20 years. These advances have served to reduce the requirements on the number of ‘particles’ necessary to faithfully represent the physics and contributed to dramatically increasing the accuracy and realism of the particle-in-cell simulation technique. An alternative approach in kinetic simulations is the Vlasov or ‘continuum’ method [10], which involves the direct solution of the kinetic equation governing the distribution function (examples include the Boltzmann and gyro-kinetic equations) on a fixed Eulerian grid in both coordinate and velocity space. Progress in the development of associated codes in recent years has also had a significant impact on the ability to realistically simulate microturbulent transport phenomena.

In dealing with macroscopic phenomena, the coupled MHD and Maxwell system of equations is first solved in the equilibrium limit, generally producing nested magnetic flux surfaces. MHD instabilities are then addressed via fluid models by advancing velocity moments of the kinetic equation in time. These include important non-ideal MHD descriptions which include resistive and viscous effects [11]. Further advances with respect to including longer mean free path dynamics are captured in so-called extended MHD models [12], 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. In general, extensive theoretical analysis over the years has led to refinements of the fluid model and improved the closure relations so that many non-fluid effects, such as particle motion and wave–particle resonances, can be represented at some level. This is illustrated in figure 4, which shows the increasingly complex levels (from left to right) of macroscopic simulation capabilities along with some representative physics applications amenable to the respective models [12].

In the rest of this section, representative examples will be highlighted of large-scale simulations that have been performed in the macroscopic stability (section 2.1), microscopic turbulent transport (section 2.2), wave–particle interactions (section 2.3) and plasma boundary physics (section 2.4) areas. Associated discussions will address possible future research directions along with the enhanced level of computational resources that will be needed.

### *2.1. Simulations of macroscopic (MHD) physics*

As noted earlier, in dealing with the space and time scales associated with large-scale instabilities in magnetically confined plasmas, it is appropriate to work with fluid-like equations, which involve solving the velocity moments of the kinetic equations coupled to the Maxwell equations. Although this is more tractable than direct kinetic simulations, many key macroscopic simulation problems in fusion science nevertheless 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. Some examples of advances of macroscopic computations which have contributed significantly to advancing understanding of toroidal plasmas include simulations of: (a) large-scale disruptions in the actual fusion-energy-producing toroidal fusion test reactor (TFTR) tokamak; (b) nonlinear evolution of magnetic islands and (c) nonlinear




**Figure 4.** Categories of macroscopic simulations with corresponding areas of physics applications.

evolution of internal reconnection events in spherical torus experiments such as the national spherical torus experiment (NSTX) in the US and the START experiment in Europe. TFTR is an example of a conventional tokamak or ‘doughnut-shaped’ torus, while NSTX and START are examples of a spherical torus with a hole through the middle, like a ‘cored apple’.

One of the most dramatic events that can occur in fusion plasmas is the major disruption—a catastrophic global collapse of a plasma discharge associated with the near-instantaneous break-up of the magnetic flux surfaces and immediate loss of confinement. The associated electromagnetic forces (e.g. from the disruption of the plasma current) and thermal loads (e.g. from the rapid quenching of the plasma temperature) can also cause major damage to the surrounding system in a tokamak device. A major research challenge for macroscopic simulations is accordingly to be able to reliably predict the conditions under which disruptions occur. This involves gaining an understanding of the physical processes responsible for triggering such events, for the nonlinear evolution after onset and for the characteristics of the final state of the various types of disruptions that have been observed. An example of some progress in this area is shown in figure 5. This depicts results from a nonlinear macroscopic simulation of a disruption event induced by high plasma pressure which prevented achievement of higher fusion power output in the TFTR tokamak [13]. The three-dimensional structure of select pressure contours shown here is associated with a complex, highly nonlinear process that involved two distinct plasma instabilities interacting to further destabilize one another nonlinearly. Basically, as an unstable long wavelength internal mode saturates, it can produce a helical structure with a local region of unfavourable curvature. Alignment of this region with the global unfavourable curvature region of the torus then produces a pressure-driven ballooning-type instability. In the presence of finite resistivity, the resultant interactions can rapidly destroy the original nested magnetic surfaces, thereby leading to a high-plasma-pressure-induced disruption. 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






**Figure 5.** Three-dimensional nonlinear macroscopic simulation of a disruption event in the TFTR tokamak experiment.


and a subsequent rapid loss of plasma current to occur. Insights gained from these studies have helped to motivate subsequent experiments as well as more realistic simulations to advance knowledge in this key area.

A fundamental and formidable challenge for MHD simulations involves dealing with the nonlinear evolution of ‘magnetic islands’. This is a true multi-scale problem in that such topological changes in the magnetic field simultaneously involve the slow (resistive) time scale associated with the growth of the islands together with the rapid Alfvén transit time scale of the underlying waves. The islands evolve over a time period that is given approximately by the ratio of the magnetic diffusion time to the Alfvén transit time (the Lundquist number,  $S$ ), i.e. a time scale in units of a normalized time for an Alfvén wave to traverse the system. When the longer mean free path dynamics resulting from toroidal geometry effects are taken into account, the resultant ‘neoclassical tearing modes’ are found to limit the maximum attainable pressure in long-duration tokamak experiments. Figure 6 illustrates the structure of the complex magnetic topology that can develop in associated nonlinear MHD simulations [14] where such plasma instabilities cause the magnetic surfaces to break into islands when the winding numbers on those surfaces are ratios of small integers. This type of simulation is presently carried out with unrealistically small values of the Lundquist number ( $S \sim 10^4$ ) to accommodate computational requirements. In addition to the major computational challenge of dealing with more realistic values (i.e.  $S \sim 10^8$  and greater for burning plasmas in fusion devices), a self-consistent closure for the neoclassical fluid equations remains under active analytic development.

As highlighted in figure 3, the development of simulation models with higher physics fidelity demands continuous validation against experimental results as well as against analytic theory and against other codes. Aided by the successful implementation of unstructured mesh algorithms and the application of advanced visualization resources to deal with complex three-dimensional toroidal structures, comparisons of advanced MHD simulations against experimental results have become increasingly compelling. For example, results from nonlinear simulations of an internal magnetic reconnection event in the NSTX and START spherical torus experiments are good illustrations of such progress. As depicted in figure 7, the sequence of images shows the evolution of an internal MHD instability in NSTX, illustrating




**Figure 6.** Resistive MHD instabilities can lead to complex magnetic topology including islands.



**Figure 7.** Simulation of the nonlinear evolution of an internal magnetic reconnection event in the NSTX.

in detail how the hot inner region of the plasma (marked in red) interchanges with the cold outer region (green) via the magnetic reconnection process [11, 15]. Comparison of these simulations with actual experimental results obtained from soft x-ray measurements of the thermal response, indicate an encouraging level of agreement. When carried out at higher resolution on more powerful computing platforms, it is expected that the level of agreement will be further improved. Similar dynamical processes are observed in the nonlinear simulation of an internal reconnection event in the START experiment and depicted in figure 8. A large



**Figure 8.** Simulation of the nonlinear evolution of an internal magnetic reconnection event in the START spherical torus experiment.

pressure bulge is seen here to grow spontaneously in a localized region on the plasma surface during the initial state and then pushes the magnetic field lines aside [16].

Although the preceding illustrative examples of advances in simulations of macroscopic dynamics have focused on toroidal (tokamak) applications, progress in modelling capabilities is also recognized to be very important for many other types of magnetic confinement configurations. For example, reversed field pinch (RFP) plasmas exhibit a nonlinear dynamo that sustains the toroidal flux against resistive diffusion. The dynamo results from the nonlinear interaction between long-wavelength, low-frequency dynamical modes. The associated magnetic fluctuations determine in large part the confinement properties of the RFP configuration. With the aid of advanced computational models, critical questions for the future of this concept will be able to be better addressed. These include multi-scale issues such as how the dynamo activity and associated magnetic fluctuations scale with the Lundquist number, and how this system behaves for times much longer than the resistive decay time of the first wall. Another prominent example is the field reversed configuration (FRC), where advanced simulation results have provided key insights into the experimentally observed behaviour. In particular, the results of recent nonlinear three-dimensional hybrid (MHD with kinetic ion dynamics) simulations offer an explanation of the stability properties observed in low- $\bar{s}$  FRC experiments, where  $\bar{s}$  is a measure of the number of ion gyro-radii in the system. The simulations using the HYM code [17] show that, although the  $n = 1$  tilt mode is linearly the most unstable mode for nearly all experimentally relevant FRC equilibria, it saturates nonlinearly without destroying the configuration, provided the FRC kinetic parameter,  $\bar{s}$ , is sufficiently small. The saturation of the  $n = 1$  tilt instability occurs in the presence of ion toroidal rotation, and is accompanied by the growth of  $n \geq 2$  rotational instabilities, which are often seen in experiments. Large- $\bar{s}$  simulations show no saturation of the tilt mode, and there is a slow nonlinear evolution of the instability after the initial fast linear growth. Overall, these hybrid simulations have provided valuable insights about the importance of nonlinear effects in determining the saturation of instabilities in low- $\bar{s}$  configurations and also for the increase in FRC life time compared to MHD models in high- $\bar{s}$  configurations [17, 18].

In looking towards the future, it is clear that significant increases in computing power and improved algorithms will be needed to accommodate the space and time resolution


demands embodied by extended MHD models and low-Lundquist-number resistive MHD investigations. Advances in computational methods, such as implicit time-stepping and adaptive mesh refinement (AMR), hold promise for reducing the computational requirements for a model of a given mathematical complexity. In addition, expected improvements in single-processor optimization, algorithmic advances and the increased availability of dedicated compute cycles should continue to stimulate impressive progress. Enhanced computational resources should also positively impact the further development of codes involving particle closures to resolve velocity space. As these computationally demanding models are exercised with sufficient spatial resolution and for longer times, perhaps a new generation of MHD ‘tools for discovery’ will emerge from the improved understanding of the physical differences between fluid and particle closures.

## 2.2. *Simulations of microturbulence*

Even if the larger-scale macroscopic disturbances in a magnetically confined plasma could be avoided, the inherent free energy (such as the expansion free energy in the temperature and density gradients) can still drive turbulent cross-field losses of heat, particles and momentum. In fact, such increased (‘anomalously large’) transport is experimentally observed to be significantly greater than levels expected from the collisional relaxation of toroidally confined plasmas (‘neoclassical theory’). This is particularly important for fusion because the effective size (and therefore cost) of an ignition experiment will be determined largely by the balance between fusion self-heating and turbulent transport losses. The growth and saturation of the associated drift-type micro-instabilities [19] have been extensively studied over the years because understanding this turbulent plasma transport process is not only an important practical problem but is generally recognized as a true scientific grand challenge. With the advent of increasingly powerful supercomputers, it is generally agreed that this problem is particularly well-suited to be addressed by modern terascale MPP computational resources.

Building on the continuous progress in this area, significantly improved models with efficient grids aligned with the magnetic field have now been developed to address realistic three-dimensional (toroidal) geometry with both global and local approaches [20]. Averaging over the fast gyro-phase of the familiar Maxwell–Boltzmann system of equations produces the gyro-kinetic equation governing the distribution function in a five-dimensional phase space [19]. As noted earlier, solution approaches that have been pursued include the particle-in-cell method, which follows the gyro-averaged orbits of an ensemble of discrete particles in a Lagrangian formulation, and the continuum (Vlasov) method, which directly solves the gyro-kinetic equation on a fixed Eulerian grid in both coordinate and velocity space. With regard to the geometry of the problems addressed, the ‘flux tube’ codes can concentrate on the fine-scale dynamical processes localized to an annular region depicted in figure 9. The associated coordinates can be described as being extended along equilibrium magnetic field lines, while being localized in the perpendicular directions. Global codes have the more imposing multi-scale challenge of capturing the physics both on the small scale of the fluctuations (micro-instabilities) and the large scale of the equilibrium profile variations. Improved implementation of gyro-kinetic particle-in-cell algorithms as well as gyro-kinetic continuum (Vlasov) approaches have been productively advanced [21].


If reliably implemented, high-resolution simulations of the fundamental equations governing turbulent transport can provide a cost-effective means to address key phenomena that would otherwise require expensive empirical exploration of a huge parameter space. The progress in capturing the ion dynamics has been impressive. For example, studies of




**Figure 9.** Flux tube simulation results of turbulence localized to an annular region of a three-dimensional toroidal plasma.

electrostatic turbulence suppression produced by self-generated zonal flows within the plasma show that the suppression of turbulence caused by prominent instabilities driven by ion temperature gradients (ITGs) is produced by a shearing action which destroys the finger-like density contours which promote increased thermal transport in a three-dimensional toroidal system [22]. This dynamical process is depicted by the sequences shown in figure 10. The lower panels, which show the nonlinear evolution of the turbulence in the absence of flow, can be compared against the upper panel sequence which illustrates the turbulence decorrelation caused by the self-generated  $\mathbf{E} \times \mathbf{B}$  flow. This is also a good example of the effective use of powerful supercomputers (e.g. the 5 teraflop IBM-SP). Typical global particle-in-cell simulations [22, 23] of this type have used one billion ( $10^9$ ) particles with 125 million ( $10^6$ ) grid-points over 7000 time-steps to produce significant physics results. In particular, large-scale simulations have been carried out to explore some of the key consequences of scaling up from present-day experimental devices (around 3 m radius for the largest existing machines) to those of reactor dimensions (about 6 m). As shown in figure 11, transport driven by electrostatic ITG turbulence in present scale devices can change character in larger systems. This transition from Bohm-like scaling to Larmor-orbit-dependent ‘gyro-Bohm’ scaling is a positive trend, because simple empirical extrapolation of the smaller system findings would be more pessimistic. Some experimental observations in a number of representative present-day experiments indicate that the relative level of turbulent heat loss increases with plasma size while the size of these eddies remains the same [24]. However, exceptions to this trend, where gyro-Bohm-like scaling that is sensitive to plasma rotation was observed, have also been reported in certain high-confinement (‘H-mode’) cases [25]. Such experiments on confinement scaling properties remain a challenging area of investigation. Nevertheless, for the larger-sized reactor-scale plasmas of the future, the present simulations would suggest that the relative level of turbulent heat loss from electrostatic turbulence does not increase with size. The underlying causes for why such a transition might occur around the 400 gyro-radii range indicated by the simulations have been explored and theoretical models based on the spreading of turbulence have been proposed [26]. Although this predicted trend is a very favourable one, the fidelity of the analysis needs to be further examined by investigating additional physics effects which






**Figure 10.** Turbulence reduction via sheared plasma flow compared with the case with flow being suppressed.



**Figure 11.** Full torus particle-in-cell gyro-kinetic simulations (GTC) of turbulent transport scaling.

might alter the present predictions. The analysis of associated scientific issues will naturally demand more comprehensive physics models within microturbulence codes. In addition to addressing experimental validation challenges, the interplay between analytic theory and advanced simulations will be increasingly important. For example, in addition to the turbulence spreading theory noted, progress in physics understanding of the nonlinear processes associated with zonal flow dynamics has resulted both from directions provided by analytic theory as well as by simulation results which have inspired new analytic models [27–30].

As noted in the preceding text, the continuum, or Vlasov, method to carry out gyrokinetic simulations of microturbulence provides a very effective complementary approach to the



**Figure 12.** Comparison of ion diffusivities from continuum gyro-kinetic simulations (GYRO) with experimentally inferred diffusivities from the DIII-D tokamak.

particle-in-cell method. In particular, the GYRO code [31], which solves the governing equations on a fixed Eulerian grid, has been able to incorporate the largest number of relevant physical processes to date. As a natural extension of the short-wavelength (‘high- $n$ ’) ballooning-type representation implemented in local flux-tube codes, this approach has effectively included passing and trapped electrons as well as electromagnetic fluctuations at finite plasma pressure. Indeed, the development of GYRO was heavily influenced by and thoroughly benchmarked against the Eulerian flux tube (local) code GS2 [32], which is actively utilized at present by the experimental community to assess confinement properties in tokamaks. Interfaced against realistic geometry and including radial profile variation with sheared  $\mathbf{E} \times \mathbf{B}$  rotation and associated parallel velocity shear (i.e. radially non-local properties), the GYRO code has successfully reproduced prominent anomalous transport trends observed in the DIII-D tokamak facility. In particular, the two DIII-D low-confinement (L-) mode discharges considered were dimensionally similar plasmas differing only in normalized gyro-radius ( $\rho_*$ ) [33] which are nearly indistinguishable stability-wise, when analysed by local, linear gyro-kinetic calculations. Figure 12 (from [34]) summarizes the calculated ion diffusivity for a total of eight simulations that were carried out in order to gauge the relative influence of various physical effects on transport observed in the experiment. The cases include (a) full physics, (b) no collisions, (c) no  $\mathbf{E} \times \mathbf{B}$  equilibrium sheared rotation and (d) zero beta. In combination, these runs suggest electron collisions act to reduce the linear ITG mode growth rate thereby reducing transport, while finite-beta effects tend to stabilize ITG modes, and simultaneously destabilize kinetic ballooning modes. Results from the most physically complete of these cases, i.e. case (a), approach the experimental results within a factor of two, with the trends from both these simulations and experiments indicating a Bohm-like scaling (diffusivity varying linearly with  $1/\rho_*$ ). The primary message from figure 12 is that as more physics is added to the simulation, the calculated diffusivity approaches that of the experiment, from being larger by a factor of 8 to falling within a factor of two, using the background profiles in temperature and density as measured in the DIII-D discharges examined. This also provides qualitative support for the importance of taking into account collisionality and  $\mathbf{E} \times \mathbf{B}$  shear effects in such calculations. It should also be noted that if the background temperature gradient is reduced by 10%, which is well within the experimental error bars, the agreement with the experimental trends becomes even

better. Since the experimental  $\mathbf{E} \times \mathbf{B}$  shearing rates are nearly independent of  $\rho$ -star, they would normally not be associated with non-gyro-Bohm transport scaling. Nevertheless, the ion diffusivity scaling is closer to gyro-Bohm when  $\mathbf{E} \times \mathbf{B}$  is neglected than when it is included. This is an illustration of the mechanisms through which gyro-Bohm scaling (intrinsic to the flux tube limit) can be broken [35]. The nature of this secondary effect is not understood at present, and highlights the importance of carrying out global (radially non-local) simulations with  $\mathbf{E} \times \mathbf{B}$  shear and collisions. Additional important ideas that might bring closer agreement are presently being explored. Although still in progress, the studies noted here indicate that previous nonlinear flux tube studies, which neglected the profile and  $\mathbf{E} \times \mathbf{B}$  shearing effects, have tended to overestimate the experimental levels of transport by roughly an order of magnitude. The computational requirement of these simulations readily challenges the latest supercomputing facilities, with a single case taking 104 processor-hours on an IBM SP-POWER3 supercomputer. Runs of this magnitude would have required roughly one year, and double the maximum memory, to complete on the largest CRAY YMP machines of the early 1990s.

An important multi-scale challenge for particle-in-cell kinetic simulations involves dealing with the realistic implementation of complete electron ('non-adiabatic') physics (including important kinetic effects, such as trapping in equilibrium magnetic wells, drift motions and wave-particle resonances) and electromagnetic dynamics. These effects have largely been incorporated into gyro-kinetic flux tube (local) codes [36], and present capabilities in gyro-kinetic global codes for dealing with electrostatic perturbations have been successfully extended to include non-adiabatic electrons [37]. Much more challenging for the global simulations are the electromagnetic perturbations, which can alter the stability properties of the electrostatic modes and also generate separate instabilities associated with deformations of magnetic surfaces. In fact, answering the long standing question about what causes the ubiquitously observed anomalously large electron thermal transport is probably linked to the ability to deal with magnetic perturbations. They can potentially cause a great increase in electron heat flux either through transient deformations of the magnetic field ('magnetic flutter') or, more plausibly, by producing an ergodic region in which the magnetic field lines no longer rest on nested flux surfaces but wander instead through a finite volume 'breaking' the flux surfaces.

In general, significant challenges for gyrokinetic simulations remain in extending present capabilities for dealing with electrostatic perturbations to include magnetic perturbations in cases where they are sufficiently large to alter the actual geometric properties of the self-consistent magnetic field. In such circumstances, micro-instabilities can drive currents parallel to the equilibrium magnetic field, which in turn produce magnetic perturbations in the perpendicular direction. These kinetic electromagnetic waves can modify the stability properties of the electrostatic modes or act as separate instabilities, such as kinetic ballooning modes for instance [38], which can alter the magnetic topology. In this sense, the kinetic simulations would encounter the major multi-scale task of also dealing with the larger-scale phenomena associated with the aforementioned MHD studies. While this fully kinetic computational challenge is yet to be substantively addressed, it should be noted here that, as also discussed within the context of extended MHD simulation models [12], nonlinear kinetic hybrid simulations involving both Vlasov-MHD and PIC-MHD models have been utilized to study the influence of energetic particles on electromagnetic waves [39].

Another major challenge for microturbulence simulations involves dealing with the highly localized regions associated with the spontaneous formation of transport barriers. These regions, in which the turbulent transport is greatly reduced or completely vanishes, are associated with regions of high electric-field shear [40] and very steep pressure gradients.




Key unresolved questions include the identification of the physics behind their radial structure and propagation, as well as accounting for their experimentally observed power threshold. Moreover, as indicated by experimental observations in these barrier regions, the electron transport can be large where the longer-wavelength modes (e.g. ITG) are predicted to be absent. This has suggested the possible relevance of the electron temperature gradient (ETG) mode, which has been actively studied [41]. The associated instabilities, which demand a kinetic description of the electron species, are again driven by expansion free energy but are predicted to be active only at extremely short wavelengths (scale-lengths between the electron gyro-radius and the collisionless skin depth). If such modes prove to be as important as the more familiar longer wavelength instabilities, dealing with the associated electron/ion mass ratio issues (e.g. simultaneously resolving ITG and ETG turbulence) would introduce a huge spatial scale problem that would be quite a formidable challenge even for terascale simulation capabilities.

At the other extreme, there are also the challenges imposed by experimentally observed phenomena which are both extremely rapid and non-local in character. These include, for example, global interactions involving the propagation of heat pulses. While rough theoretical models based on pictures of avalanches and self-organized criticality have been proposed, detailed simulation investigations still await systematic exploration. While much work remains before a realistic computational model that captures all of the key physics behind microturbulence-driven transport is at hand, encouraging progress has been made. It is encouraging that present flux-surface-averaged models of heat transport, developed by fitting the ion thermal conductivity to results from nonlinear microturbulence simulations, have achieved reasonable agreement when compared with magnetic confinement experimental results [42].

### *2.3. Simulations of wave–particle interactions: RF and energetic particle physics*

Naturally occurring or laboratory produced plasmas can support a variety of electromagnetic waves, which play a fundamental role in the dynamics of the plasmas either in the form of externally driven waves or as self-generated instabilities. Arguably the most prominent type is the Alfvén wave, discovered by Nobel laureate Hannes Alfvén. Electromagnetic waves connect the fluid-like motion of the continuum background with the discrete particle-like motion of the individual charged particles through intricate resonances, such as Landau and cyclotron resonances. This connection opens up the possibility of transferring energy and momentum between the waves and the plasma, which has led to robust techniques for increasing the temperature of laboratory plasmas to the fusion-relevant regime. Such experimentally validated methods have also provided the non-inductive current needed for steady-state fusion reactor operation. The vital role of computations in this key area of research is to provide realistic quantitative calculations of particle motions in background electromagnetic fields. Accurate results from these computations are needed to assess the efficacy of various applications of waves to heat a particular species of the plasma and also to address the impact of energetic particles which result from strong auxiliary heating and/or the generation of alpha particles from fusion reactions.

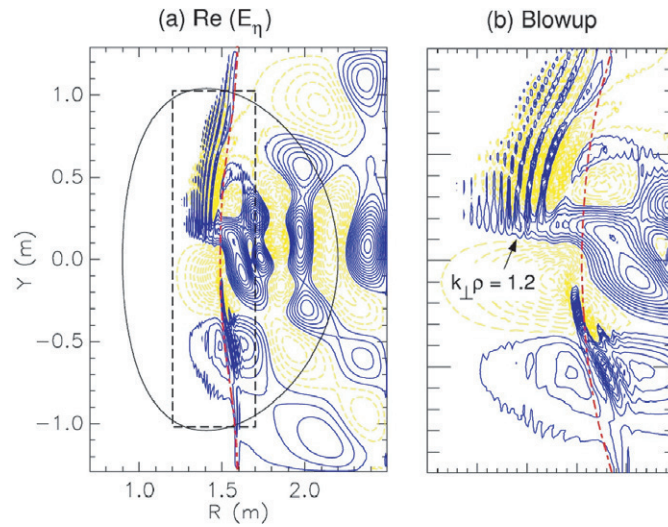
To respond to the challenge just described, simulation codes of wave–plasma interactions are being developed which make increasingly effective use of terascale computers. Impressive progress has been made in the massive parallelization and acceleration of computer-intensive full-wave radiofrequency (RF) field solvers, extension of all-order methods to two- and three-dimensional plasmas, increasing and incorporating details needed to accurately account for the wave–particle resonances in the RF conductivity models and benchmarking of code results




**Figure 13.** Reduction in spurious cyclotron damping with improved poloidal mode resolution in TORIC simulation results for D ( $^3\text{He}$ ) mode conversion experiments in the Alcator C-Mod tokamak.

against experimental measurements. An illustrative recent example is the application of MPP resources to resolve short-wavelength structures in mode conversion studies at the ion cyclotron range of frequency (ICRF). In practice, long-wavelength electromagnetic waves are launched into the plasma externally via a complicated antenna structure. As the waves propagate toward the high-temperature core of the plasma, they enter into regions where a variety of internal plasma waves can co-exist. Depending on plasma parameters, the externally launched ‘fast’ wave can interact strongly with one or more of these internal waves and transfer energy to them in the process. The internal waves can, in turn, effectively transfer their energy to the charged particles via wave–particle resonances. Calculating the multiple-branching process of wave propagation and energy transfer with high accuracy has been a great challenge for RF heating research [43].

Rapid progress has recently been made in improving the performance of full-wave, two-dimensional ICRF solvers, such as TORIC [44], via advances in the algorithm for matrix inversion and adaptation to a massively parallel platform (MPP) [45, 46]. The improvement has eliminated spurious damping that has previously masked the transfer of wave energy from the ‘fast’ wave to the ion Bernstein wave (IBW), which is in turn absorbed via electron Landau damping. This behaviour, which is illustrated in figure 13, is in good agreement with experiment [46]. Complementary to the TORIC code, which is limited to relatively long wavelengths and cyclotron harmonics of two or less by the small gyro-radius expansion approximation, an alternate full-wave code (AORSA) has been developed [47]. The AORSA code removes these restrictions by using a fully spectral representation for the wave-field, a Cartesian coordinate system, and a co-location method of discretization. This new capability also allows self-consistent modelling of high harmonic ‘fast’ wave heating under study in the low-aspect-ratio torus experiment, NSTX. However, in order to apply this new method to two- and three-dimensional problems, intensive demands are placed on computer speed and memory, which can only be satisfied by accessing high-performance computers. As shown in figure 14, such calculations have been used, for example, to produce high-definition solutions for mode conversion in two-dimensional geometry. Results from both the TORIC and AORSA codes indicate that, for high minority ion concentrations, mode conversion is dominated by transfer of wave energy from the ‘fast’ wave to a slow ion cyclotron wave (ICW) instead of an IBW. This is consistent with earlier simplified analytic models [48] and experimental measurements [49]. The ICW propagates in the opposite direction to the IBW and can be damped by both the ions and the electrons. Finally, as illustrated in figure 15, the AORSA



**Figure 14.** Two-dimensional field solutions for conversion of fast ion cyclotron waves to ion Bernstein waves in the heating of the DIII-D tokamak plasma.



**Figure 15.** All-orders spectral calculation of minority ion cyclotron heating for all ten field periods of the large helical device (LHD) stellarator facility with a single antenna located at the extreme right-hand side. Individual cross-sections show the logarithm of the minority ion power absorption at various toroidal angles.

code has been extended to give fully three-dimensional solutions for minority ion heating in stellarator geometry [47, 50]. This problem is significantly more challenging than the two-dimensional problem, and improvements will be needed to efficiently apply the same methods to three-dimensional physics studies.


The preceding examples illustrate the challenges in modelling wave heating in fusion plasmas which are natural consequences of the multiple branch possibilities for wave–wave and wave–particle interactions. High accuracy and resolution are required to correctly model

the linear wave propagation dynamics. Since the present computational solutions of the wave equations already involve 120 000 coupled complex linear equations for  $200 \times 200$  Fourier modes, future demands for more processors and memory will be severe. In addition, the need to simulate the nonlinear evolution of such phenomena poses further formidable challenges. These requirements may be significantly alleviated in the future with more efficient advanced algorithms, such as adaptive mesh refinement approaches, and the development of a more efficient formulation, moving from the current dense matrix form of the linear wave problem to a sparse matrix form. Scientific progress in this area will be important for future development of wave-based techniques for plasma control needed to help access and then control plasma behaviour in advanced confinement regimes. From an integrated modelling perspective, which will be discussed later in this paper, it will also be important to address the question of how RF heating will impact the MHD, transport and edge physics properties of a magnetically confined plasma. As a final note, it should be pointed out that in addition to fusion applications, advances in these computational wave models will also be of interest to areas such as (a) RF-driven plasma sources for plasma processing of materials including semiconductors [51], (b) advanced space-propulsion approaches [52]; and (c) possible relevance of ion cyclotron damping of Alfvén waves to the high-temperature solar wind [53].

#### *2.4. Simulations of plasma boundary physics*

Plasma boundary dynamics, which includes plasma–wall interactions and turbulent transport in the geometrically-complex plasma edge region, represents a most formidable challenge from both the physics formulation and the computational modelling perspectives. Improved capabilities for simulating interactions between the plasma and its material interfaces have largely come into being in the tokamak divertor research area. Key activities here include the development of models dealing with atomic physics and the presence of neutral gas which is able to penetrate into the plasma edge. Proper formulation of the basic equations describing complex experimentally observed processes such as filamentation and bursting behaviour are also needed to advance the physics understanding that is needed. In addition, present laminar models need to be generalized to more realistic three-dimensional descriptions. This will require strong interplay between development of advanced simulation codes and systematically benchmarking them against theoretical models motivated by experimentally observed trends. In particular, it is necessary to realistically simulate the behaviour of the large fraction of the hot plasma from the interior as it is neutralized near the divertor surface and pumped away. The modelling challenge here also includes complex effects such as the heating of this divertor surface and the conversion of heat energy to radiation. With regard to atomic physics at the plasma edge, advanced computational studies have addressed a wide range of atomic collisional processes present in controlled fusion plasmas. This has involved the implementation of state-of-the-art atomic and molecular collision codes on terascale computing facilities. Accurate atomic data is essential for development of a realistic collisional radiative model and new results have been obtained on associated topics, including electron-impact ionization of lithium and beryllium, dielectronic recombination and ion–atom collisions [54].


In order to achieve the much needed advances in the understanding of transport at the plasma boundary, a better physics description of edge turbulence is clearly needed. Though related to aspects of the core turbulence issues discussed earlier in this review, this topic is even more challenging in many ways, including dealing with the absence of the usual scale-separation paradigms. In a magnetically confined plasma, the last closed flux 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 the open ones that strike a material surface.



**Figure 16.** Wave number ( $k$ ) spectrum comparisons between gas-puff imaging (GPI) experimental measurements and three-dimensional electromagnetic fluid code simulations of edge turbulence on the Alcator C-Mod tokamak.

Along the open field lines, the plasma particles, momentum and heat are rapidly conducted to the material surface. A boundary layer is accordingly formed about the last closed flux surface where the radial gradients steepen until the cross-field transport can compete with the rapid transport parallel to the equilibrium magnetic field. Analysis of this boundary layer is greatly complicated by the aforementioned absence of applicable scale separation orderings. In particular, the ion orbit width, which provides the characteristic scale of the turbulence, is actually comparable in magnitude to the scale length of the equilibrium density and temperature gradients. Some advances in understanding have nevertheless been gained from applications of electromagnetic fluid codes to simulate the edge plasma turbulence [55]. For example, figure 16 shows a favourable level of agreement in the comparisons of the wave number ( $k$ ) spectrum from three-dimensional electromagnetic fluid code simulations of edge turbulence with results from gas-puff imaging experimental measurements on the Alcator C-Mod tokamak [56]. This example, which should encourage much needed further studies, also illustrates the kind of substantive physics validation noted in section 2 of this review. The associated challenge for experiments is to produce and implement advanced experimental diagnostics which are capable of producing integrated measurements of key physical quantities such as spectra, correlation functions, heating rates and other variables of interest.

Overall, research in the area of edge turbulence and modelling studies will need to address a number of the key scientific issues, which include: (a) the self-consistent coupling of core heat and particle fluxes with edge sources; (b) the effects of the separatrix, the scrape-off-layer (SOL), and the  $x$ -point geometry in toroidal plasmas; (c) moving beyond fluid equations to deal with kinetic dynamics, which involve addressing a variety of compatibility issues associated with various turbulence models, such as gyro-kinetic ions and multi-fluid electrons, and full kinetic electrons and ions; (d) dealing with nonlinear structures and intermittency; (e) gaining an understanding of the transition to and the development of the L to H transition; (f) assessing the role of atomic and molecular physics in the structure and dynamics of the edge plasmas and (g) taking into account the physical effects of proximity to MHD stability boundaries in more realistic assessments of edge localized mode (ELM) phenomena. It is of particular interest to gain an understanding of this last topic because ELMs have been observed in the edge thermal barrier region characteristic of discharges with the best energy confinement. Both large and small amplitude ELMs apparently act to reduce/regulate the temperature pedestal height when the thermal gradient becomes too steep.




**Figure 17.** Three-dimensional gyro-kinetic global particle-in-cell codes have demonstrated excellent scaling as the number of processors is increased.

### 3. Terascale computations

In order to effectively deal with the challenging scientific issues highlighted in this review, the plasma science community must address advanced code development tasks which are important for most areas of research. The basic goal of enhancing the physics fidelity of the codes and developing the significantly improved software to deal with highly complex problems involves addressing: (a) multi-scale physics such as kinetic electromagnetic dynamics which have been discussed in earlier sections of this review; (b) more efficient algorithms compatible with evolving computational architectures and (c) scalability of codes necessary for utilizing terascale platforms. As the computational hardware advances to meet the demands from the largest, most difficult problems, it is essential to also meet the continuing challenge of improving the scientific applications software and the associated algorithms. Powerful approaches, such as adaptive mesh refinement for higher dimensionality phase space, will need to be actively pursued.

With regard to efficiently implementing present generation codes on the most powerful MPP computers, it is encouraging that for virtually all of the topical areas covered in this review, the plasma science community has had success in developing codes for which computer run-time and problem-size scale well with the number of processors. A good example of this trend is illustrated in figure 17, where the global microturbulence PIC code, GTC, has demonstrated excellent scalability for more than 2000 processors on the IBM-SP computer at the National Energy Research Supercomputer Center (NERSC). This code is the representative from the fusion energy science area within the NERSC suite of demonstration/benchmark codes to evaluate realistic performance on new advanced computational platforms. Active collaboration on the world's most powerful supercomputer, the Earth Simulator Computer (ESC) in Japan, has just commenced and involved the recent porting of this code to the ESC site. The goal of this ongoing project is to evaluate the importance of the ESC's vector-parallel architecture in comparison to the much more widespread super-scalar MPP architecture. An active collaboration has also been initiated with a complementary global particle simulation effort in Japan [57]. Results from these early benchmark runs were quite impressive. Specifically, utilization of 64 ESC processors yielded results which were not only



**Figure 18.** Terabytes of data are now generated at remote locations, such as the heat potential shown here on 121 million grid points from a particle-in-cell turbulence simulation.

more efficient (by about a factor of two) but were more than 20% faster than 1024 processors on the IBM-SP supercomputer at NERSC. Efficiency in this context refers to measured ability of a given code to achieve the theoretically rated performance level of the computer processors. In addition to the ESC, the new X1 vector supercomputer at the Oak Ridge National Laboratory has commenced benchmarking activities involving both the continuum code, GYRO [31], as well as the particle-in-cell code, GTC [58]. Overall, the practical goal here is to effectively utilize the tools, technologies and advanced hardware systems that will help minimize the time-to-solution for the most challenging computational plasma physics problems.

It should be emphasized that a natural consequence of the effective utilization of supercomputers is the tremendous amount of data generated, as illustrated in figure 18. Terabytes of data are even now generated at remote locations (e.g. where supercomputing centres are located), presenting data management and data grid technology challenges [59]. The data must be efficiently analysed to compute derived quantities. New advanced visualization techniques are needed to help identify key features in the data. There are also significant programming and algorithmic challenges, which must be met in order to enable computational capabilities for addressing more complex scientific problems. These include multi-dimensional domain decomposition in toroidal geometry and mixed distributed/shared memory programming. Other problems include load balancing on computers with large numbers of processors, optimization of fundamental gather-scatter operation in particle-in-cell codes and scalable parallel input/output (I/O) operations for the petascale range of data sets.

Another key priority to help accelerate progress on the impressive state-of-the-art physics advances involves developing a set of diagnostic and visualization tools that will allow real-time interaction with the simulated data. This will impact the ability of scientists to effectively test theories/hypotheses and to address specific questions about the proper implementation of key physics within the computational models. Recent efforts using GKV, an interactive data analysis and visualization tool expressly designed for analysing output from plasma microturbulence simulations, provide a good example of significant advances in this area [20]. Also, in order to realize the benefits from advancements in understanding, it will be necessary to periodically



update existing integrated models to ensure that they reflect the fresh insights gained from these new ‘tools for discovery’. This point is further elaborated upon in the next section of this review, dealing with the subject of integrated modelling challenges in fusion energy science.

The development of diagnostic instruments capable of making high-resolution measurements of electric and magnetic fields and cross-sectional measurements of turbulent fluctuations has made it increasingly feasible to demonstrate more in-depth correlations between experimental results and theoretical models. This has the potential to greatly improve the basic understanding of the mechanisms controlling plasma confinement. As in dealing with the output from terascale simulations, maximizing the effectiveness of such simulation/experiment comparisons will also necessitate addressing critical computer science and enabling technology (CSET) issues in the area of data management and visualization. Effective utilization of the power of advanced computing to solve challenging problems can be best exploited when the necessary infrastructure is established and effective software tools are made available. Terascale computing requires complementary software that scales as well as the hardware and which provides an efficient code development environment. In general, improved networking, data management and visualization are needed to strengthen the coupling of terascale simulations with theory and experiment.

Since the challenges in applications infrastructure development for terascale computing usually involves multiple research institutions, system integration is needed, together with the availability of software that allows maximal use of available computing platforms. Modern object-oriented code development methods are accordingly required to facilitate sharing code development efforts among collaborators from numerous research groups. A good example of recent progress in the plasma physics area involves direct collaborations with the CSET community. As emphasized throughout the course of this review, it is critically important for simulations to be rigorously validated against experiments to ensure the fidelity of the science. Advances in this direction will require improvements in connectivity to experimental data and to state-of-the-art tools for data visualization, mining and manipulation. An illustrative example is provided by the partnership formed between computational applications scientists at the three largest US magnetic fusion experimental facilities at General Atomics (San Diego, CA), the Princeton Plasma Physics Laboratory (Princeton, NJ) and the Massachusetts Institute of Technology (Cambridge, MA) with computer scientists at Argonne National Laboratory Lawrence Berkeley National Laboratory, and the Computer Science Departments at the University of Utah and Princeton University. This cross-disciplinary effort [60] has made encouraging progress toward goals which include: (a) more efficient use of experimental facilities via powerful new between-pulse data analysis capabilities; (b) better access by researchers to analysis and simulation codes, data and visualization tools and (c) creation of a standard tool set for remote data access, security and visualization. Future international collaborations on large-scale burning plasma experiments such as International Thermonuclear Experimental Reactor (ITER) will demand enhanced capabilities of this kind.

As a final point, it should be kept in mind that even with access to greatly improved computational hardware and software advances, there will remain limitations to what can be practically achieved [61]. Indeed, some of the most complex plasma phenomena involving highly transient nonlinear behaviour may defy mathematical formulation and be beyond the reach of computational physics.

#### **4. Integrated modelling challenges in fusion energy science**

In many respects, the physics fidelity of the most advanced high-performance computational codes, when properly benchmarked against theory, experimental results and complementary



codes, represents the state of understanding in any research discipline. The integrated modelling challenge in fusion energy science is to effectively harvest the knowledge gained from the associated simulations to provide the underpinning for predicting the behaviour of fusion systems. Developing a comprehensive simulation capability for carrying out ‘virtual experiments’ on such systems will be essential for the design and optimization of a portfolio of future facilities, including burning plasma experiments, technology testing facilities and demonstration power plants, necessary for the realization of commercially available fusion energy. A more near-term application is to optimize the design of the first experimental burning plasma device, which will pave the way to both greater scientific understanding and speeding up the development of the first fusion power plant. A realistic integrated simulation capability would dramatically enhance the utilization of such a facility, and its application to concept innovation can be expected to lead to further optimization of toroidal fusion plasma systems in general. The targeted goal is to deliver the ability to effectively simulate the entire plasma device with a validated predictive capability that reproduces trends from existing experimental regimes and can be applied with a reasonable level of confidence when extrapolated to new physics regimes.

To reiterate some key points made earlier in this review, the characteristics of fusion plasmas make the goal of integrated modelling an extremely challenging one. Because fusion plasmas are naturally subject to both large- and small-scale disturbances, which relax the system to a lower energy state, representing the associated physics involves multiple time scales, ranging over fourteen orders of magnitude and multiple spatial scales, ranging over eight orders of magnitude (figure 1). In some cases, such as those in the edge plasma region, the large- and small-scale phenomena are strongly coupled, thereby invalidating conventional separation of scale treatments. Moreover, the computational domains are geometrically complex, the solutions severely anisotropic, and the underlying physical processes are coupled with essential nonlinearities. In short, the algebraic systems that must be solved are often ill-conditioned, and the physics approximations currently invoked are not completely justified. As a consequence, the desired set of governing simulation equations do not yet rest on a firm physics foundation. Nevertheless, there have been promising approaches that have been developed for each of these challenges in an individual sense [62, 63]. However, a proper integrated simulation capability for fusion plasmas will need to effectively respond to all of these challenges simultaneously.

One comprehensive approach toward the systematic achievement of the integrated modelling goal in fusion energy science involves demonstrating significant advances in three phases. At the most fundamental level, improvements in physics understanding and theoretical descriptions for all physical processes in key areas that govern the performance of fusion systems are needed. This should translate into a capability to perform realistic numerical simulations of individual components of a fusion device, utilizing available high performance computers to help quantitatively validate them against accurate experimental measurements. For example, in the US, this initial step has been successfully launched under the US Department of Energy’s SciDAC program [2]. A second level is the optimization of physics code packages and demonstration of coupled simulations of several different physics processes. This element will require significantly greater collaborations with the mathematics and computer science community for improvement of algorithm accuracy and efficiency along with the preparation of physics packages for compatibility with the next generation of high-performance computer architectures. The associated advanced simulation capability should have extensive ability to better diagnose and interpret experiments. The third and final level is to integrate all physical processes needed in a seamless framework for the comprehensive computational simulation of fusion energy science experimental devices. Achieving this impressive integrated simulation goal in a timely manner (guided by the three-level phased

approach just described) would demand a major increase in funding support for researchers dedicated to such a mission. It would also require terascale (and eventually petascale) ‘capability’ computing resources as well as significantly enhanced ‘capacity’ computing capabilities, possibly available via dedicated topical computing centres.

A complementary approach to addressing the challenge of integration is to focus on improving the physics modules and modernizing the computational framework of existing one- (or ‘one-and-a-half-’) dimensional predictive integrated transport codes. While the associated tool which emerges here is not likely to be as comprehensive and revolutionary as the first approach noted, it is also likely to be a more immediately usable new modelling capability. Nevertheless, capturing the key elements of more realistic higher dimensionality simulations into simplified modules remains a formidable task. Significant efforts in this area are being initiated in Europe and Japan as well as in the United States.

## 5. Cross-disciplinary alliances: progress and opportunities

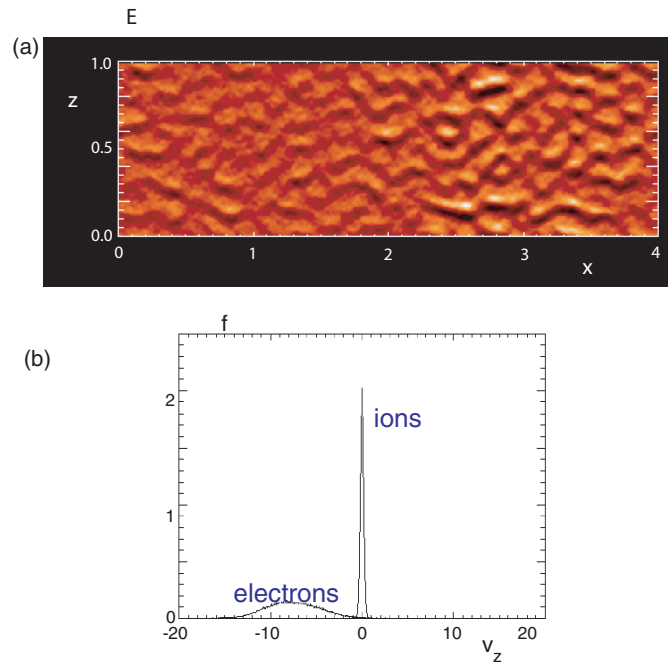
The increased emphasis on advanced computations in fusion energy sciences has helped provide opportunities for attracting and assimilating the bright young talent needed for the future health of this field. It has also served to stimulate mutually beneficial scientific alliances with other applications areas. For example, even though they reside in very different parametric regimes, the effective modelling of global systems (climate), combustion and fusion devices all deal with complex, three-dimensional, nonlinear fluid flows and associated kinetic dynamics. They share the common computational challenge of rapidly developing advanced integrated modelling capabilities that are capable of treating complex dynamical systems covering many decades in time and space. The computational challenges posed by nonlinear plasma fluid problems also share many common features with computational fluid dynamics (CFD) issues faced in global systems modelling, materials sciences and numerous other areas. Long-standing plasma physics programmes in the space physics and basic science arenas provide excellent opportunities, inviting further strengthening. In addition, advanced computational modelling capabilities will also stimulate progress in non-fusion plasma applications, including the manufacture of plasma display panels, plasma thrusters for satellites, microelectronic components 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, space physics, solar physics and astrophysics. Opportunities currently exist not only for productive cross-disciplinary interactions between scientific application areas with applied math and computer science, but also for cross-cutting collaborations among various application disciplines themselves. Illustrative examples are described in the following.

(1) Magnetic reconnection simulations: Magnetic reconnection is a fundamental dynamical process which enables a magnetized plasma system to convert magnetic energy into high-speed flows, thermal energy and energetic particles. This is a research topic of great interest in plasma physics as well as other areas such as space physics and astrophysics [64]. The multi-scale nature of this class of problems comes from the challenge of understanding the kinetic processes occurring at very small scales which control the release of magnetic energy on a global scale. Specifically, it is necessary to account for both the dynamics in narrow layer regions where long-mean-free-path kinetic dynamics are important and the distant regions where a macroscopic description with proper boundary conditions enforced is appropriate. The associated research for this multi-scale grand challenge problem, which involves the application of both MHD and kinetic analysis capabilities, has produced significant theoretical advances [65]. This has been complemented by dedicated laboratory experiments which have actively explored the structure of the current sheet which defines the ‘dissipation region’ [66].


Computational progress in magnetic reconnection studies in laboratory plasmas provides a good example of how plasma physics advances can productively impact other areas of research. Particularly noteworthy is the influence on the exploration of associated phenomena in magnetospheric and solar physics research. Such findings can be relevant to various aspects of solar and space physics, including topics such as solar flares, coronal mass ejection and magnetic sub-storms in the magnetosphere. A good example comes from the collaboration led by the University of Maryland and involving researchers at the Max Planck Institute for Plasma Physics, the University of Minnesota and Dartmouth College. This has involved particle simulation applications using 70 M particles and 20 M grid points in three-dimensional simulations with periodic boundary conditions. Results here suggest that the large reduction of the magnetic reconnection current layer width is associated with the rapid acceleration of electrons [7]. As illustrated in figure 19, the resultant ‘electron beam’ is found to produce a two-stream instability that nonlinearly evolves into ‘electron holes’—localized regions of intense anti-parallel electric field where the electron density is depleted. The authors have pointed out that the simulated structure of these electron holes compares favourably with satellite observations at the Earth’s magnetopause. They point out that the birth and death of these electron holes and their associated intense electric fields can lead to strong electron scattering and energization, whose understanding is critical to explaining why magnetic explosions in space release energy so quickly and produce such a large number of energetic electrons.

(2) Design of advanced experiments: When applied to the design of highly complex systems, such as, the National Compact Stellarator Experiment (NCSX) [68], the efficient utilization of massively parallel processing (MPP) computations proved to be essential for the optimization of the key stability and transport properties. The quasi-axisymmetric stellarator depicted in figure 20 is governed by a large number of physical mechanisms and processes that are described by many variables. Since the construction and operation of experimental facilities to advance the requisite knowledge base can be slow and costly, new science-based design tools are needed to improve upon such ‘cut and try’ methods. The challenge here involved applying new integrated simulation capabilities made possible by modern supercomputers. In addition to efficiently combining the key physics properties from the best current theoretical models, the MPP capabilities enabled extensive campaigns of computer simulations on a huge number of designs in a short time and at low cost. New optimization software was specifically developed to allow a systematic exercising of all the relevant variables to home in on the optimal design. In contrast to tokamaks, whose shape can be described with only four variables, it takes about 40 variables to account for the more complicated shapes of stellarators. The potential advantages of stellarators include efficient sustainability in the steady state and being less prone to sudden plasma termination events (i.e. major disruptions). However, the design permutations which needed to be explored for this more complex device accordingly approached nearly a million variations. Quickly finding the most promising designs in such a large landscape posed a formidable challenge in computational complexity. Nevertheless, enabled by access to modern MPP dedicated ‘capacity’ computing resources, integrated simulations, which combined physics models with practical engineering criteria, successfully produced the NCSX design. Specifically, the MPP computations were critical for addressing ‘constructability’ issues which involved the proper design of the magnetic coil system that could produce the plasma of interest. This is an excellent example of how simulations enabled by modern supercomputers can provide an attractive cost-effective new path for moving from innovative research concepts to the design of practical experiments.

(3) Modelling of advanced diagnostics: To gain a better understanding of the complex processes impacting plasma turbulence, theorists, computer scientists and experimentalists have recently developed an MPP simulation of the actual microwave reflectometry




**Figure 19.** Three-dimensional particle-in-cell code simulations of dynamics in the magnetic reconnection current layer leading to the generation of localized regions of intense anti-parallel electric field ('electron holes').



**Figure 20.** Massively parallel supercomputers effectively utilized to design future experimental facilities via optimization of confinement and constructability properties.

diagnostic [69]. As illustrated in figure 21, the new MPP code models the actual incoming microwaves reflecting off a target plasma which evolve as its turbulence grows. This is interfaced with the latest large-scale microturbulence simulation results described earlier [22, 23]. The characteristic turbulent correlation length can be extracted from this reflectometer




**Figure 21.** Interfacing of MPP simulation of microwave reflectometry diagnostic with large-scale microturbulence simulation results.


simulation and compared with that deduced from the original microturbulence simulation. This innovative new capability can cost-effectively aid both in the interpretation of the reflectometry data from experiments as well as helping to more efficiently utilize the existing diagnostic capabilities and to design improved diagnostic capabilities.

(4) High-energy and accelerator physics: The areas of high-energy and accelerator physics provide another good example of how computationally driven advances in plasma physics can aid progress in other disciplines. In particular, a particle simulation approach (‘delta-f’ method) pioneered in magnetic confinement studies [9], has been incorporated into a three-dimensional particle code for simulating both ions and electrons in an accelerator environment. This code has successfully simulated electron–proton two-stream instabilities from initial noise [70]. As shown in figure 22, results have been obtained from tracking about  $10^6$  particles over  $10^6$  lattice periods (betatron oscillations). This has addressed a key problem of interest to the high energy and accelerator community—the presence of an unwanted electron population which has been observed in the Proton Storage Ring Experiment at the Los Alamos National Laboratory and is of possible relevance to the interpretation of behaviour in the Spallation Neutron Source Experiment at the Oak Ridge National Laboratory. The associated physics is also of interest in heavy ion beam fusion research [71]. These new advanced MPP simulation studies provided the interesting proposition that this potentially troublesome instability could be eliminated with modest axial momentum spread.

As a final comment, it is important to emphasize that the future health of the plasma physics research area will depend on how successfully it can find better ways to attract, educate and then assimilate bright young talent into the field. Computational plasma science provides an excellent path forward to address this challenge. Indeed, many of the key advances cited in this review have involved major contributions from young scientists. Nevertheless, greater efforts are required to actively develop and nurture opportunities for training the next generation of researchers with capabilities that cross-cut traditional disciplinary boundaries. This is true in other areas of science as well and has motivated a number of major universities in the US to respond by establishing interdisciplinary graduate educational programmes that target the development of capabilities that ‘bridge’ various areas of computational science applications



**Figure 22.** Three-dimensional particle-in-cell simulation of electron–proton two-stream instability.



**Figure 23.** Example of the kind of graduate programme which promotes cross-disciplinary education in various areas of computational science applications (including plasma physics) together with applied math and computer science.

(including plasma physics) together with computer science and applied mathematics. An illustrative example is provided by the ‘PICASso’ programme at Princeton University—a graduate programme in integrative computer and applications science. Figure 23 shows schematically that collaborating departments and participants foster interdisciplinary research for the mutual benefit of all participating areas.

## 6. Conclusion

The ‘computational grand challenge’ nature of plasma physics in general and fusion research in particular is a consequence of the fact that in addition to dealing with vast ranges in space and time scales which can span over ten decades, the relevant problems involve extreme anisotropy, the interaction between large-scale fluid-like (macroscopic) physics and fine-scale kinetic (microscopic) physics and the need to account for geometric detail. Moreover, the requirement of causality (inability to parallelize over time) makes this problem among the most challenging in computational physics. There has been excellent progress

during the past decade in fundamental understanding of key individual phenomena in high-temperature plasmas. 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 termination of the discharge. Major advances have been achieved in the modelling of such dynamics, which require an integration of fluid and kinetic physics in complex magnetic geometry. Significant progress has also been made in addressing the dynamics governing the breaking and reconnection of magnetic field lines, which is a central scientific issue for fusion energy as well as allied fields such as astrophysics and space and solar physics. Another key topic, where there have been exciting advances in understanding, is the degradation of confinement of energy and particles in fusion plasmas caused by turbulence associated with small-spatial-scale plasma instabilities that are driven by gradients in the plasma pressure. While progress has been impressive, the detailed physics of the growth and saturation of these instabilities, their impact on plasma confinement and the knowledge of how such turbulence might be controlled remain major scientific challenges.

Accelerated development of computational tools and techniques is vitally needed to develop predictive models which can prove superior to empirical scaling. For example, this will have a major impact on the fusion research community's ability to effectively harvest the key physics from the proposed ITER. The probability that ITER will achieve its goals can be significantly enhanced by development of the capability to numerically simulate the plasma behaviour under realistic conditions. The challenge to unravel the mystery of the complex behaviour of strongly nonlinear, non-equilibrium plasma systems, including interactions with their external environments, is clearly the next frontier of computational fusion research and will advance the understanding of fusion plasmas to an exciting new level. This will be made possible by the dramatic advances in supercomputer technology which will enable simulations of increasingly complex phenomena with greater fidelity.

The scientific issues of magnetic fusion encompass a wide range of disciplines including those just mentioned, as well as others. However, the dynamics of high-temperature plasmas do not respect these categorizations, and the understanding of overall plasma performance requires integrating all of these issues in an integrated simulation that includes interactions between phenomena which were previously studied as essentially separate disciplines. To achieve the ultimate goal of such integration, it becomes necessary to follow the evolution of the global profiles of plasma temperature, density, current and magnetic field on the energy-confinement time scale with the inclusion of relevant physics on all important time scales. While this is a formidable long-term goal, the field now stands ready to begin such cross-disciplinary studies and to increase the physics content of existing integrated codes. This will necessitate the development of an architecture for bringing together the disparate physics models, combined with the algorithms and computational infrastructure that enables the models to work together. The associated integration of code modules, many of which are individually at the limits of computational resources, will clearly require substantial increases in computer power. With the rapid advances in high-end computing power, researchers in plasma physics as well as other fields can expect to be able to model such systems in far greater detail and complexity, leading eventually to the ability to couple individual models into an integrated capability which can enable significantly improved understanding of an entire system.

Supercomputing resources can clearly accelerate scientific research critical to progress in plasma science in general and to fusion research in particular. Such capabilities are needed to enable scientific understanding and to cost-effectively augment experimentation by allowing efficient design and interpretation of expensive new experimental devices (in the multi-billion dollar range). In entering the exciting new physics parameter regimes required to study burning plasmas, the associated challenges include higher spatial resolution, dimensionless

parameters characteristic of higher temperature plasmas, longer simulation times and higher model dimensionality. It will also be necessary to begin integrating these models together to treat nonlinear interactions of different phenomena. Various estimates indicate that increases by factors of  $10^3$  to  $10^5$  are needed in the combined computational power. Associated challenges include advancing computer technology, algorithmic development and improved theoretical formulation—all contributing to the improvement of overall ‘time-to-solution’ capabilities.

Overall, the exciting advances in information technology and scientific computing have enabled key contributions to all areas of plasma science and the acceleration of progress toward gaining the physics knowledge needed to harness fusion energy. The associated achievements and approaches highlighted in this review hold great promise for improving scientific understanding of experimental data, for stimulating new theoretical ideas and for helping produce innovations leading to the most attractive and viable designs for future facilities. This has helped transform traditional research approaches and has provided a natural bridge for fruitful collaborations between scientific disciplines leading to mutual benefit to all areas. As a natural consequence, computational plasma science also provides an excellent path forward to address the challenge of attracting, educating and retaining the bright young talent essential for the future health of the field.

## Acknowledgments

The authors are grateful for the excellent input material provided by the many members of the plasma physics advanced computational community and especially to Dr Greg Rewoldt for his assistance in helping us to integrate this information. Their impressive productivity is evident in the examples cited and has provided the basis for the increased awareness and appreciation by the general scientific community of the high quality of research and the stimulating challenges in the area of computational plasma physics. This work was supported by US Department of Energy contracts DE-ACO2-76CH03073 and DE-AC03-89ER51114.

## References

- [1] Tang W M 2002 *Phys. Plasmas* **9** 1856
- [2] Department of Energy, Office of Science 2001 *Scientific Discovery Through Advanced Computing Program (SciDAC)* [www.science.doe.gov/scidac/](http://www.science.doe.gov/scidac/)  
Dunning T 2001 private communication
- [3] National Research Council, Fusion Science Assessment Committee 2001 An assessment of the Department of Energy’s Office of Fusion Energy Sciences program Final Report (Washington, DC: National Academy Press)
- [4] Rutherford P H and Frieman E A 1968 *Phys. Fluids* **11** 569  
Taylor J B and Hastie R J 1968 *Plasma Phys.* **10** 479  
Catto P J, Tang W M and Baldwin D E 1981 *Plasma Phys.* **23** 639  
Frieman E A and Chen L 1982 *Phys. Fluids* **25** 502
- [5] Freidberg J P 1987 *Ideal Magnetohydrodynamics* (New York: Plenum)  
White R B 2001 *Theory of Toroidally Confined Plasmas* (London: Imperial College Press)
- [6] Dawson J M 1983 *Rev. Mod. Phys.* **55** 403  
Cohen B I, Barnes D C and Dawson J M 1995 *Comp. Phys. Commun.* **87** 1
- [7] Langdon A B and Birdsall C K 1970 *Phys. Fluids* **13** 2115
- [8] Lee W W 1983 *Phys. Fluids* **26** 556  
Lee W W 1987 *J. Comput. Phys.* **72** 243
- [9] Dimits A M and Lee W W 1993 *J. Comput. Phys.* **107** 309  
Parker S E and Lee W W 1993 *Phys. Fluids B* **5** 77



- [10] Knorr G 1962 *Nucl. Fusion* **3** 1119  
Cheng C Z 1977 *J. Comput. Phys.* **24** 348  
Denavit J and Kruer W L 1971 *Phys. Fluids* **14** 1782
- [11] Wesson J A 1976 *Comput. Phys. Commun.* **12** 53  
Rosenbluth M N *et al* 1976 *Phys. Fluids* **19** 1987  
Biskamp D 1997 *Nonlinear Magnetohydrodynamics* (New York: Cambridge University Press)
- [12] Park W *et al* 1999 *Phys. Plasmas* **6** 1796  
US Department of Energy (DOE) *Fusion SciDAC Extended Magnetohydrodynamic Modeling Project* 2001 at <http://w3.pppl.gov/CEMM/>
- [13] Park W *et al* 1995 *Phys. Rev. Lett.* **75** 1763
- [14] Glasser A H *et al* 1999 *Plasma Phys. Control. Fusion* **41** A747  
Sovinec C R *et al* 2001 *Phys. Plasmas* **8** 475
- [15] Park W *et al* 2003 *Nucl. Fusion* **43** 483
- [16] Hayashi T, Mizuguchi N, Watanabe T, Todo Y, Sato T and the Complexity Simulation Group, NIFS 2000 *Nucl. Fusion* **40** 721  
Mizuguchi N *et al* 2000 *Phys. Plasmas* **7** 940  
Hayashi T *et al* 2002 *Proc. 19th IAEA Fusion Energy Conf. 2002 (Lyon, France)* paper TH/6-3
- [17] Belova E V, Davidson R C, Ji H and Yamada M 2004 Kinetic effects on the stability properties of field-reversed configurations: II. Nonlinear evolution *Phys. Plasmas* **11** 2523
- [18] Belova E V, Davidson R C, Ji H and Yamada M 2003 *Phys. Plasmas* **10** 2361
- [19] Tang W M 1977 *Nucl. Fusion* **18** 1089  
Horton W 1999 *Rev. Mod. Phys.* **71** 735
- [20] DOE Fusion SciDAC Plasma Microturbulence Project 2001 <http://fusion/gat/com/theory/pmp>  
Nevins W M 2001 private communication
- [21] Waltz R E, Candy J and Rosenbluth M N 2002 *Proc. 19th IAEA Fusion Energy Conf. 2002 (Lyon, France)* paper TH/P1-19
- [22] Lin Z *et al* 2002 *Phys. Rev. Lett.* **88** 195004
- [23] Lin Z *et al* 2002 *Proc. 19th IAEA Fusion Energy Conf. 2002 (Lyon, France)* paper TH/1-1
- [24] Budny R V *et al* 2000 *Phys. Plasmas* **7** 5038  
Rhodes T L 2001 APS-DPP Invited Paper UI1-5 *Bull. Am. Phys. Soc.* **46** 323
- [25] Petty C C *et al* 2002 *Phys. Plasmas* **9** 128
- [26] Hahn T S, Diamond P H, Lin Z, Itoh K and Itoh S-I 2004 *Plasma Phys. Control. Fusion* **46** A323
- [27] Rosenbluth M N and Hinton F L 1998 *Phys. Rev. Lett.* **80** 724
- [28] Chen L, Lin Z and White R 2000 *Phys. Plasmas* **7** 3129
- [29] Diamond P H *et al* 2001 *Nucl. Fusion* **41** 1067
- [30] Malkov M A, Diamond P H and Rosenbluth M N 2001 *Phys. Plasmas* **8** 5073
- [31] Candy J and Waltz R E 2003 *J. Comput. Phys.* **186** 545
- [32] Dorland W *et al* 2000 *Proc. 18th IAEA Fusion Energy Conf. 2000 (Sorrento, Italy)* paper TH2/5
- [33] McKee G R *et al* 2001 *Nucl. Fusion* **41** 1235
- [34] Candy J and Waltz R E 2003 *Phys. Rev. Lett.* **91** 045001
- [35] Waltz R E, Candy J and Rosenbluth M N 2002 *Phys. Plasmas* **9** 1938
- [36] Chen Y and Parker S E 2001 *Phys. Plasmas* **8** 2095  
Chen Y and Parker S E 2003 *J. Comput. Phys.* **189** 463
- [37] Lin Z and Chen L 2001 *Phys. Plasmas* **8** 1447  
Lee W W *et al* 2001 *Phys. Plasmas* **8** 4435
- [38] Tang W M *et al* 1980 *Nucl. Fusion* **20** 1439
- [39] Todo Y *et al* 1995 *Phys. Plasmas* **2** 2711  
Todo Y and Sato T 1998 *Phys. Plasmas* **5** 1321
- [40] Burrell K H 1997 *Phys. Plasmas* **4** 1499  
Kishimoto Y *et al* 1999 *Plasma Phys. Control. Fusion* **41** A663  
Parail V V 2002 *Plasma Phys. Control. Fusion* **44** A63
- [41] Jenko F *et al* 2000 *Phys. Plasmas* **7** 1904  
Dorland W *et al* 2000 *Phys. Rev. Lett.* **85** 5579  
Jenko F and Dorland W 2002 *Phys. Rev. Lett.* **89** 225001  
Labit B and Ottaviani M 2003 *Phys. Plasmas* **10** 126
- [42] Dimits A M *et al* 2000 *Phys. Plasmas* **7** 969
- [43] Itoh K, Itoh S-I and Fukuyama A 1984 *Nucl. Fusion* **24** 13  
Villard L 1986 *Comput. Phys. Rep.* **4** 95

- [44] Brambilla M and Krücken T 1988 *Nucl. Fusion* **28** 1813  
Brambilla M 1999 *Plasma Phys. Control. Fusion* **41** 1
- [45] Wright J C *et al* 2004 *Phys. Plasmas* **11** 2473
- [46] Bonoli P T *et al* 2000 *Proc. 18th IAEA Fusion Energy Conf. 2000 (Sorrento, Italy)* paper EXP4/01
- [47] Batchelor D B *et al* 2002 *Proc. 19th IAEA Fusion Energy Conf. 2002 (Lyon, France)* paper TH/P3-21  
DOE Fusion SciDAC Wave–Plasma Interactions Project 2001 <http://www.ornl.gov/fed/scidacr/>
- [48] Perkins F W 1977 *Nucl. Fusion* **17** 1197
- [49] Nelson-Melby E *et al* 2003 *Phys. Rev. Lett.* **90** 155004
- [50] Jaeger E F *et al* 2001 *Phys. Plasmas* **8** 1573  
Jaeger E F *et al* 2002 *Phys. Plasmas* **9** 1873
- [51] Lieberman M *et al* 1994 Design of high-density sources for materials processing *Plasma Sources for Thin Film Deposition and Etching* ed M H Francombe and J L Vossen (San Diego, CA: Academic)
- [52] Chang-Diaz F 2000 The VASIMR rocket *Sci. Am.* **283** 90
- [53] Cranmer S R 2000 *Astrophys. J.* **532** 1197
- [54] Loch S D *et al* 2004 *Phys. Rev. E* **69** 066405  
Pindzola M S, Minami T and Schultz D R 2003 *Phys. Rev. A* **68** 013404
- [55] Hallatschek K and Zeiler A 2000 *Phys. Plasmas* **7** 2554  
Rogers B *et al* 1998 *Phys. Rev. Lett.* **81** 4396  
Xu X Q *et al* 2000 *Phys. Plasmas* **7** 1951  
Jenko F and Scott B D 1999 *Phys. Plasmas* **6** 2705
- [56] Zweben S J *et al* 2002 *Phys. Plasmas* **9** 1981  
Zweben S J *et al* 2004 *Nucl. Fusion* **44** 134
- [57] Idomura Y and Tokuda S 2003 *Nucl. Fusion* **43** 234
- [58] Oliker L *et al* 2003 Evaluation of cache-based superscalar and cacheless vector architectures for scientific computations *Proc. of the Supercomputing (SC) 2003 Conf. (Phoenix, Arizona, US, Nov. 2003)* <http://www.sc-conference.org/sc2003/paperpdfs/pap255.pdf>  
Ethier S 2003 private communication
- [59] Klasky S *et al* 2003 Grid-based parallel data streaming implemented for the gyrokinetic toroidal code *Proc. of the Supercomputing (SC) 2003 Conf. (Phoenix, Arizona, US, Nov. 2003)* <http://www.sconference.org/sc2003/paperpdfs/pap207.pdf>
- [60] DOE Fusion SciDAC Collaboratory Project 2003 [www.science.doe.gov/scidac/](http://www.science.doe.gov/scidac/)  
Schissel D P for the Fusion Collaboratory Team 2003 Building the US national fusion grid: results from the National Fusion Collaboratory Project *Proc. of the 4th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research (San Diego, CA, July 21–23, 2003)* *Fusion Engineering and Design* **71** 245
- [61] See for example Laughlin R B 2002 Physical basis of computability *Comput. Sci. Eng.* **4** 27
- [62] Garbet X 2003 *Nucl. Fusion* **43** 975
- [63] Joffrin E *et al* 2002 *Nucl. Fusion* **42** 235
- [64] Vasilunas V M 1975 *Rev. Geophys. Space Phys.* **13** 303  
Parker E N 1979 *Cosmical Magnetic Fields* (Oxford: Clarendon)  
Parker E N 1973 *Astrophys. J.* **180** 247  
Sato T and Hayashi T 1979 *Phys. Fluids* **22** 1189  
Biskamp D 1986 *Phys. Fluids* **29** 1520  
Kulsrud R M 1998 *Phys. Plasmas* **5** 1599  
Tang W M 1998 *Science* **279** 1488
- [65] Birn J *et al* 2001 *J. Geophys. Res.* **106** 3715  
Aydemir A Y 1999 *Phys. Fluids B* **4** 3469  
Rogers B N *et al* 2001 *Phys. Rev. Lett.* **87** 195004
- [66] Yamada M *et al* 2000 *Phys. Plasmas* **7** 1781  
Yamada M 2001 *Earth Planets Space* **53** 509
- [67] Drake J F, Swisdak M, Cattell C, Say M A, Rogers B N and Zeiler A 2003 *Science* **299** 873
- [68] Reiman A *et al* 2001 *Phys. Plasmas* **8** 2083  
Hudson S R, Monticello D A, Reiman A H, Boozer A H, Strickler D J, Hirshman S P and Zarnstorff M C 2002 *Phys. Rev. Lett.* **89** 275003
- [69] Kramer G J, Nazikian R and Valeo E J 2004 *Plasma Phys. Control. Fusion* **46** 695
- [70] Qin H *et al* 2000 *Phys. Rev. ST Accel. Beams* **3** 084401  
Qin H *et al* 2000 *Phys. Rev. ST Accel. Beams* **3** 109901
- [71] Bangerter R O 1999 *Phil. Trans. R. Soc. A* **357** 575