Taylor & Francis
Taylor & Francis Group

Fusion Science and Technology

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/ufst20

How the Exascale Computing Project and Private Magnetic Fusion Research Stimulated Each Other

T. Tajima, S. A. Dettrick, Z. Lin, R. E. Groenewald, B. S. Nicks, A. Veksler, D. C. Barnes, F. Ceccherini, J. Drobny, L. Galeotti, S. Gupta, C. K. Lau, A. Necas & M. Onofri

To cite this article: T. Tajima, S. A. Dettrick, Z. Lin, R. E. Groenewald, B. S. Nicks, A. Veksler, D. C. Barnes, F. Ceccherini, J. Drobny, L. Galeotti, S. Gupta, C. K. Lau, A. Necas & M. Onofri (22 Oct 2024): How the Exascale Computing Project and Private Magnetic Fusion Research Stimulated Each Other, Fusion Science and Technology, DOI: 10.1080/15361055.2024.2397620

To link to this article: https://doi.org/10.1080/15361055.2024.2397620

FUSION SCIENCE AND TECHNOLOGY

DOI: https://doi.org/10.1080/15361055.2024.2397620

How the Exascale Computing Project and Private Magnetic Fusion Research Stimulated Each Other

T. Tajima,[a,b] S. A. Dettrick,[a]* Z. Lin,[b] R. E. Groenewald,[a] B. S. Nicks,[a] A. Veksler,[a] D. C. Barnes,[a] F. Ceccherini,[a] J. Drobny,[a] L. Galeotti,[a] S. Gupta,[a] C. K. Lau,[a] A. Necas,[a] and M. Onofri[a]

[a]*TAE Technologies, Inc., 19631 Pauling, Foothill Ranch, California 92610*
[b]*University of California, Irvine, Irvine, California 92697*

Received April 15, 2024
Accepted for Publication August 19, 2024

Abstract — *We provide a brief summary and comments on how we utilized the Exascale Computing Project (ECP) supercomputing resources for the development of the frontal research of nuclear fusion reactor development. In turn, our demand and use helped stimulate the ECP hardware and software, and our project to tackle real-world simulation of fusion experiments has stimulated the ECP and its development. We posit that thus the legacy of the ECP is important and useful and that it will continue to have a growing impact on scientific progress.*

Keywords — *Public-private partnership,* exascale*, simulation, fusion, field-reversed configuration.*

Note — *Some figures may be in color only in the electronic version.*

I. INTRODUCTION

TAE Technologies, Inc. (TAE) has been a member of the Exascale Computing Project[1,2] (ECP) Industry Agency Council (IAC) since its inception. ECP, initiated in 2016 across multiple administrations, is a collaborative venture involving government agencies and the private sector and led by the U.S. Department of Energy (DOE). The IAC was initially an Industrial Council, including participation from several major U.S. corporations in the aerospace and energy sectors and a few small companies such as TAE. Later, government agencies such as the National Aeronautics and Space Administration, National Oceanic and Atmospheric Administration, and National Institutes of Health were added to form the IAC. The ECP represents the pinnacle of high-performance computing, and TAE's role within it stems from its need for access to the world's fastest computers and accompanying software. In turn, the ECP obtains experience and evaluation of its technology chain by a tangible application (in this case, demanding high-performance fusion reactor computation). This access is vital for conducting extensive simulations related to fusion reactor development, encompassing not only plasma behavior but also various reactor components such as first walls and external source terms. Through its utilization of ECP-based computations, TAE showcases to the public and private sectors the immense potential of high-performance computing and aligns with the ECP's visionary goals.

It may be worthwhile to look back at a bit of the prehistory of supercomputing before we reached the ECP. It is noted that from the early stage of U.S. supercomputers such as CDC 6600-7600 and Cray-1 and Cray-2 in the 1970s through the 1980s, fusion research played an important role to drive supercomputing (and driving supercomputers). For example, Dawson, at the University of California, Los Angeles, was a leader of large-scale particle-in-cell (PIC) computation and its fusion plasma research methodology. He advocated and helped enhance

---

*E-mail: sean@tae.com

supercomputer centers such as the National Magnetic Fusion Energy Computer Center and National Energy Research Scientific Computing Center (NERSC), which may be regarded as an early bird of the present-day ECP. One of us (TT) worked under him to realize such computational use (see Refs. [3] and [4]). Later on, the use of supercomputers to investigate and promote fusion spawned out a sophisticated computational paradigm of fusion research via a large array of codes utilizing supercomputers; an example may be the Numerical Tokamak Project[5,6] of the mid-1990s, a community-wide initiative in collaboration with the DOE aimed at simulating fusion plasma behavior. [In the present day, two of the leaders of the Numerical Tokamak Project, Tajima and Barnes, are leaders in field-reversed configuration (FRC) fusion plasma research at TAE.] TAE's development of the Numerical Fusion Reactor (NFR) goes beyond simulating just the plasma; it aspires to simulate the entire fusion reactor, making exascale computing an essential requirement. TAE's unique position drives the boundaries of exascale computing and provides substantial motivation for advancements in ECP's hardware and software.

TAE is committed to rapidly achieving safe, commercially viable fusion reactors for global use. To fulfill this mission, TAE relies heavily on top-tier computing resources provided by DOE computing centers to expedite its goals. TAE was also privileged to be an early user of ECP state-of-the art applications, significantly enhancing its progress. Simultaneously, TAE's participation in ECP technologies allows it to assess and explore the highest levels of its own technological capabilities.

We will elucidate these two aspects of our experience since we joined the ECP IAC. In Sec. II, we briefly summarize the state of fusion research at TAE. In Sec. III, we describe our long-term efforts to build an integrated modeling suite for FRC plasmas. In Sec. IV, we describe our recent pivot to adopt ECP tools, particularly WarpX and AMReX, in our modeling suite. In Sec. V, we briefly summarize another public-private partnership (PPP), between TAE and the University of California, Irvine (UCI), involving the Gyrokinetic Toroidal Code (GTC), which also received support from the ECP. Finally, we conclude and make observations about the future ecosystem of ECP tools.

II. TAE FUSION RESEARCH

Previously known as Tri Alpha Energy, TAE was founded in 1998 with the goal to develop a novel nuclear fusion technology that is simple enough to be practically commercializable. Building upon decades of publicly funded fusion advancements and original work conceived at UCI, TAE has developed a unique approach to fusion that combines the FRC confinement concept with high-power neutral beam injection.

Our vision[7,8] is to use the p-11B fuel cycle ($p + {}^{11}B \rightarrow 3{}^{4}He$) because it has inexpensive, nonradioactive, and superabundant reactants and yields only helium. With no radioactive products from primary reactions, it should be possible to build a reactor with existing nuclear-qualified materials. Our reactor research and development use the FRC, which has long been recognized by the DOE[9,10] as a promising route to advanced fuels. The FRC is compact, with high power density and low synchrotron radiation due to the high-beta plasma, and it has axisymmetric geometry with simple circular confinement coils. We believe that the p-11B fuel cycle and FRC geometry provide a combination of regulatory advantages and engineering simplicity that make commercialization more feasible. To our knowledge, TAE is the only laboratory working on steady-state, high-beta plasmas with the intent to burn advanced fuels. In fact, we recently made the first-ever measurements of p-^{11}B fusion in a magnetically confined plasma.[11] Our research is therefore an important complement to tokamak research and could be viewed as a technological hedge in the worldwide fusion effort.

Over the past two decades, TAE has dedicated itself to refining the concept of the beam-driven FRC in the C-1,[7] C-2,[12,13] C-2U,[14] and C-2W[15,16] (also known as Norman) series of experimental devices. Notably, TAE's latest experimental device, nicknamed Norman in honor of the late Norman Rostoker, the company's founder, has not only met but also surpassed its major performance milestones. It has achieved steady-state beam-driven FRC plasmas characterized by high-beta plasma (approximately 85%), substantial fast ion pressure (3 keV total temperature), high electron temperatures (T_e exceeding 400 eV), sustained FRC plasma for 30 ms, and the successful demonstration of magnetic field ramp-up.[16]

The C-2W device has served as a test bed for advanced artificial intelligence applications, including experimental optimization,[17] real-time plasma control,[18] Bayesian inference of plasma states,[19,20] advance prediction of kinetic instabilities with neural networks,[21] and Bayesian reconstruction of radial mode shape evolution.[22]

The valuable lessons gleaned from this experimental journey have laid the foundation for TAE's next-step Advanced FRC experiment, named Copernicus. This reactor-scale prototype will use protium ^{1}H isotope plasmas to demonstrate plasma density, temperature, and confinement times that would make net-energy production in deuterium-tritium isotope fuels feasible. TAE will then construct a fusion power plant (FPP) prototype called Da Vinci to demonstrate the net-energy gain of p-^{11}B fuel.

III. A NUMERICAL FUSION REACTOR

Beyond experimentation, TAE is a leader in algorithm development and scientific code modeling.[23] Since 2002, we have been developing the NFR, a comprehensive suite of tools for the computer simulation of the FRC. These tools are instrumental in assessing reactor design scenarios, thereby mitigating risks before physical construction. The models not only encompass the entire fusion plasma but also include its surroundings. As such, NFR is far more ambitious than the earlier version of the Numerical Tokamak Project in the past.[3]

The NFR is a Grand Challenge problem in computational physics. It is a multiple timescale, multiple space scale, multiphysics, and engineering problem. Timescales in the fusion plasma span 14 orders of magnitude, from 10^{-12} s for electrostatic oscillations to 10^{-6} s for Alfvénic oscillations to 100 s for confinement and fusion reaction times. Space scales span from submillimeter electron shielding length to ~10 cm gradient scale lengths to ~10-m machine sizes.

To overcome the multiscale nature of the problem, multiple models with different levels of fidelity are developed. The physics models encompass core plasma dynamics, plasma-wall interactions, and external components. The core plasma models include fully kinetic ions, electrostatic microturbulence,[24–30] macrostability,[31–33] null magnetic field regions, magnetic mirror physics, and edge plasma physics. They connect with plasma-facing walls and external actuators through dedicated models of plasma-material interaction,[34] neutral beams,[35] radio-frequency (RF) heating,[36] plasma fueling, electrode biasing, real-time magnetic control, resistive walls, and power supplies.[37]

Integration of these various models is performed by either file-based data transfer or in-memory coupling. When infrequent information transfer between models is required, for example, when an equilibrium model is used to initialize an initial value code, file-based data transfer is performed using a standardized (within TAE) hierarchical structure in the HDF5 format[38] that describes all required aspects of the boundaries, domain, computational mesh, electric and magnetic fields and potentials, and plasma fluid quantities. Kinetic quantities (from PIC) are stored in separate HDF5 files as necessary. External actuators including coils, circuits, and neutral beams and RF heating are described in standardized formats also. When frequent information transfer between models is required, such as tight coupling between magnetohydrodynamic and kinetic models to obtain a global model of thermal plasma transport under the influence of neutral beams,[39–41] in-memory coupling between models is performed.

As python has become the dominant data analysis language in the last decade, we have also adopted python interfaces to our simulation suite. Since the C-2W device geometry and every external actuator and diagnostic are completely described in the Machine State Database[42] (MSDB) and all shot-to-shot settings and diagnostic results are stored in the MDSplus[43,44] database, both of which have python interfaces, the complete set of inputs and outputs on C-2W are machine readable and amenable to automated analysis. Because of the complexity of the system, we eschew the development and maintenance of graphical user interface tools in favor of python kernel Jupyter notebooks, which can be used for the complete workflow of reading the machine state from MSDB and MDSplus, configuring simulation code inputs, job dispatch through the queuing system, analysis, and visualization. For large three-dimensional (3D) visualizations, we use Paraview[45] or Visit.[46]

Utilizing these models and connectivity to experiment, TAE has conducted high-fidelity predictive simulations of FRC fusion plasma experiments.[47] These simulations have harnessed the computational might of leadership-class computing resources at the DOE's NERSC, Argonne Leadership Computing Facility (ALCF), and Oak Ridge Leadership Computing Facility (OLCF), through awards of computing time TAE received under the auspices of DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE), ASCR [Advanced Scientific Computing Research] Leadership Computing Challenge (ALCC), and Energy Research Computing Allocations Process (ERCAP) programs.

IV. TAE'S INTEGRATION OF EXASCALE RESOURCES INTO FUSION RESEARCH

The Exascale[a] Computing Initiative, a partnership between the DOE Office of Science (SC) and National

---

[a] An Exascale computer is one that has the capability to perform 10^{18} floating point operations per second.

Nuclear Security Administration, was a 7-year project, launched in 2016. One major component was the delivery of new exascale-class computers at DOE national laboratories, and another component was the ECP. The ECP mission is to ensure all the necessary pieces are in place for the nation's first, capable, exascale ecosystem, including critical applications and an integrated software stack. It has been instrumental in creating public software tools designed to empower scientific simulation applications, enabling them to harness the immense computational potential of the world's fastest supercomputers.

Up until 2022, TAE's paradigm for high-performance-computing (HPC) code development was to write bespoke C++ and Fortran code to target distributed memory systems, using established libraries for parallel applied mathematics and computer science, including OpenMP and message passing interface (MPI) libraries, with parallel linear algebra and input/output (I/O) implemented using the PETSc,[48] Parallel HDF5,[38] and ADIOS2[49] libraries. Using these tools, we successfully demonstrated compute capability for two of our HPC codes, showing strong scaling on up to one-third of the ALCF Theta computer.

However, in mid-2022, TAE strategically pivoted to refactor its HPC algorithms within the ECP application WarpX[b] and the ECP framework AMReX.[50] This change in paradigm reduces our need for bespoke code development and allows us to leverage higher-level features of the ECP software stack that provide physics features such as particle push and PIC operations (in WarpX) on top of applied math features such as meshing, differential operators, embedded boundaries, and domain decomposition (in AMReX). Along with these higher-level features come the ECP computer science advances including optimization of instructions on chips and accelerators, shared memory and distributed memory parallelism, load balancing, resiliency, program control, parallel I/O, in situ analysis, and visualization. This transformation enables us to develop a flexible and modular computational framework with a single codebase adaptable to various HPC hardware choices.

TAE is giving back to the community and helping to bolster the legacy of the ECP by contributing its WarpX code modifications upstream to the public git repository.[51] We are adding new physics models[34,47,52] to the public repository and expanding the scope of WarpX to encompass a wide array of magnetic confinement fusion plasma scenarios. Potential applications of our new models extend beyond fusion research into space plasmas and plasma arcs.

The WarpX code implements a variety of field solvers including electrostatic, magnetostatic, and electromagnetic and has a rich set of multiphysics modules, including embedded boundaries, Monte Carlo modules for Coulomb collisions, fusion reactions, and quantum electrodynamic processes. This makes it very versatile, and by adding new models, we have further extended the scope of WarpX to include physics relevant to magnetic fusion energy (MFE) including global stability, kinetic ion instabilities, plasma-material interaction, and turbulent transport. WarpX was originally developed to simulate plasma accelerators, either laser-driven plasma accelerators or particle beam–driven accelerators, both of which use wakefield[53] acceleration.

An example fusion application of the preexisting WarpX models has been the study of the Alfvén ion cyclotron (AIC) instability.[54] This mode can increase axial loss of ions in mirrorlike geometries,[55] including the FRC. WarpX has been used to study a one-dimensional plasma composed of kinetic electrons and kinetic, anisotropic ions. The results of the simulations suggest a viable experimental strategy to explore the intentional stabilization of the AIC mode through sufficiently short plasmas.[56] See Fig. 1 for visualization of magnetic field perturbations due to the AIC mode.

For global stability studies of fusion plasmas, TAE has contributed a 3D hybrid PIC model with fully kinetic ions and massless, isotropic, neutralizing fluid electrons. Such a model was originally developed for magnetospheric plasmas[57] but has also been very useful for FRC plasmas[58–60] because it captures the electromagnetic wave-particle interactions on Alfvénic and ion inertial timescales that determine global stability of FRC plasmas. TAE has refactored its in-house MPI/OpenMP implementation of this model, FPIC,[31,32] using WarpX and AMReX. Validations of this hybrid PIC algorithm, and computational performance, have been presented in a recent publication.[47] One of the validations that has been performed with the hybrid PIC model is the simulation of magnetic reconnection in an initially force-free current sheet.

Four snapshots in time of a reconnection simulation are presented in Fig. 2. The measured reconnection rate in this simulation was found to match the value reported in the isothermal electron case in the literature.[61] The hybrid PIC model has also been validated against the growth rate of the tilt instability in fully 3D simulations. Three snapshots in time of an example tilt simulation are shown in Fig. 3.

---

[b] WarpX won the 2022 Gordon Bell prize for first-principles simulations of laser-based electron accelerators, and it was the first ECP application project to run on the full scale of Frontier and complete its ECP milestone goal.

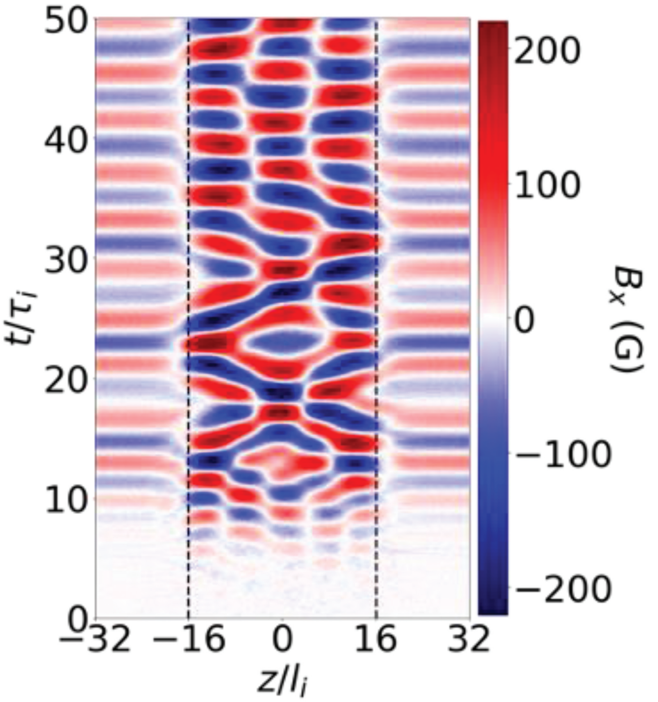

Fig. 1. One-dimensional WarpX kinetic simulation[56] of AIC instability, showing one Cartesian component of shear magnetic fluctuations. The vertical axis is normalized to the ion cyclotron period τ_i, and the horizontal axis is normalized to the ion skin depth l_i.

Typical performance of the WarpX implementation of our hybrid PIC algorithm for a reactor-relevant FRC plasma is illustrated in Fig. 4. In this study, the mesh resolution of the 3D simulation is determined by the ion inertial length (which must be resolved in hybrid PIC) and is set as 148 × 148 × 592 cells. The number of particles in each simulation is determined by available graphics processing unit (GPU) memory. We scale the number of A100 GPUs from 4 to 64 to accommodate from 60 to 960 particles/cell. This is a weak scaling study, which is more relevant than a strong scaling study when the mesh resolution is predetermined by the physics.

Figure 4 shows that we can push 0.8 to 12 billion particles in 0.5 s or less depending on the GPU count, with good weak scaling. A similar scaling plot for an identical study was shown by Groenewald et al.,[47] but the one shown here is approximately 1.5 times faster for every run. This highlights an important benefit of the collaboration. While TAE's physics test suite here has remained the same, Lawrence Berkeley National Laboratory has made important computer science updates in the background including adding GPU-aware MPI, mixed precision, and particle sort improvements, leading to a significant speedup.

TAE is also continuing to develop a Darwin PIC model[52] that we intend to add to the WarpX suite of solvers. This will be used to improve our model of electrode biasing of the FRC plasmas. Part of this new model is a semi-implicit scheme to solve the Poisson equation, already adopted in the Aleph code,[63] which alleviates the need to resolve the electron plasma frequency. This is important for long timescale simulating of the sheath and presheath at the electrode surfaces. Electrode biasing modifies radial force balance, which causes the plasma to respond on ion inertial (Alfvénic) timescales. These timescales will be captured by the magnetostatic component of the Darwin solver, with time variation carried by the PIC particle population.

Plasma-material interactions (PMIs) are of crucial importance for fusion as they determine some of the boundary conditions on plasma confinement. We have connected the open source kinetic PMI model RustBCA[34] to WarpX via the pywarpx interface to integrate the physics of reflection, sputtering, and implantation into our framework of models. We also extended the existing WarpX binary collision module, which captures Coulomb collisions and fusion reactions, to include charge exchange interactions for the modeling of neutral beam injection.

Energy confinement in MFE plasmas is to a certain extent determined by a turbulent inverse energy cascade driven by electrostatic kinetic microturbulence. Models of this process must resolve the timescale and space scale of electron dynamics, which are very short in comparison to the ion scales required for global stability, so that Leadership Computing Facility (LCF)–class supercomputers are required. A small number of nonlinear, fully kinetic, electrostatic simulations of the Vlasov-Poisson system were previously performed using TAE's OpenMP/MPI code ANC[24,28,30] and validated against experimental fluctuation measurements.[26] We are now adopting WarpX for these simulations, and the resultant increase in simulation capability on multi-GPU LCF machines will permit detailed parameter scans that will allow us to understand empirically observed trends in plasma transport in FRC plasmas.

We have also used the WarpX code for first-principles simulations of terahertz radiation generation to support the development of a new experimental measurement of the internal magnetic field of the plasma.[64] Pulsed polarimetry, a Lidar technique, requires terahertz radiation with certain qualities, and WarpX simulations show that the requirements can be met by irradiating solid targets with intense laser light.

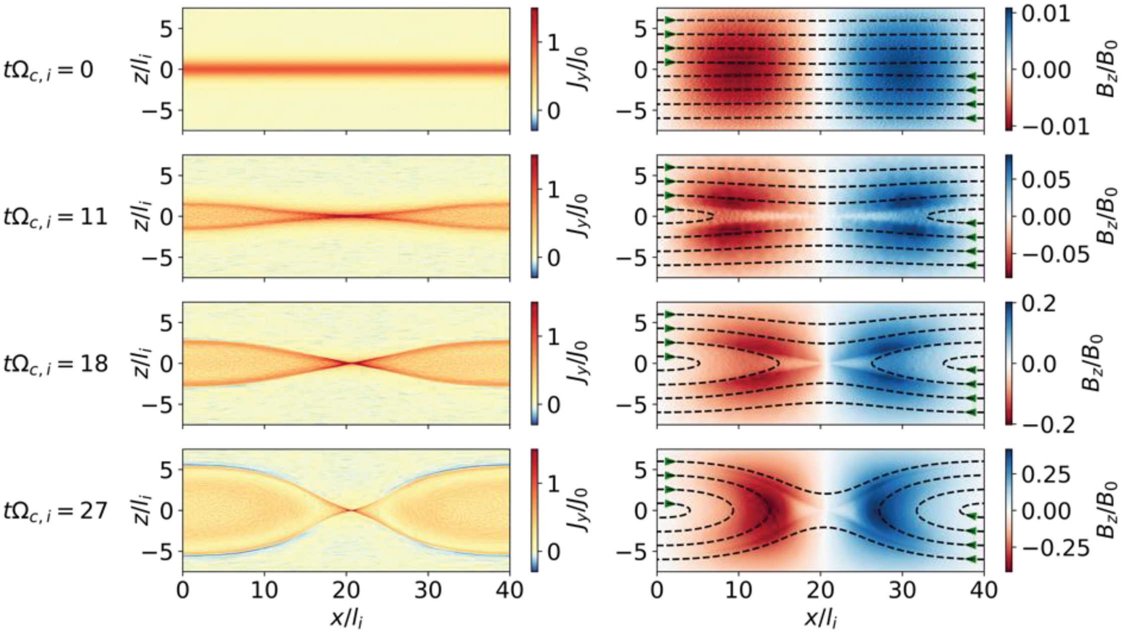

Fig. 2. Two-dimensional fluid/kinetic simulation of magnetic reconnection, seeded with a single X-point, using the new hybrid solver. Vertical and horizontal axes (z and x directions) are normalized to the ion skin depth l_i in each plot. The four rows of plots show snapshots at times $t\Omega_{ci}$=0, 11, 18, and 27, where Ω_{ci} is the ion gyrofrequency. The left column of plots shows the normalized current density in the y direction (out of plane). The right column of plots shows the magnetic field lines as dashed lines, and the z component of the magnetic field in color.

The use of WarpX and AMReX fits well into the integrated modeling workflow mentioned in Sec. III because AMReX has a python application programming interface (API), pyAMReX, which enables our newer WarpX models to couple with our existing physics codes and with the experimental databases MSDB and MDSplus. Much of our WarpX development code is written in python using pyAMReX, even the CUDA accelerated portions of the code, and because pyAMReX allows zero-copy read/write access to the GPU memory that is managed by the C++ WarpX and AMReX application codes, there is little to no performance penalty.

V. TAE AND ITS PPP PARTNERS

In addition to our participation in the WarpX community, TAE is at the center of multiple PPPs in fusion plasma physics, funded both by direct investment from TAE and by DOE instruments such as the Leadership Computing, Innovation Network for Fusion Energy (INFUSE), and Scientific Discovery through Advanced Computing (SciDAC) programs.

We are forming a computational platform that can be used to study a broad scope of plasma dynamics, algorithms, and an array of different confinement strategies. This approach facilitates validation of in-house numerical tools through high-level comparison among various approaches. Our UCI partners can be more tuned toward the tokamak and stellarator strategy while we at TAE are spearheading the FRC-based strategy. Such a stereoscopic high level of computational strategy seldom appears, unless the present PPP platform exists.

UCI's interest centers on tokamaks/stellarators and their implications for fusion reactors. Specifically, they are concerned with turbulent transport[65] and energetic particle (EP) confinement[66] and their cross-scale interaction,[67] which are some of the crucial issues for burning plasma experiments like ITER and FPPs. Ignition in these scenarios depends on self-heating by energetic fusion products, particularly alpha particles. The balance between the heating power and turbulent transport determines the performance of these fusion reactors.

There are no experimental data on alpha-particle transport in the ignition regime dominated by alpha-particle heating to extrapolate from existing fusion devices to the ITER and FPP. Therefore, predictive simulation incorporating multiscale and multiphysics is critically important for the ITER operational scenario development and the FPP design, performance assessment, and optimization.

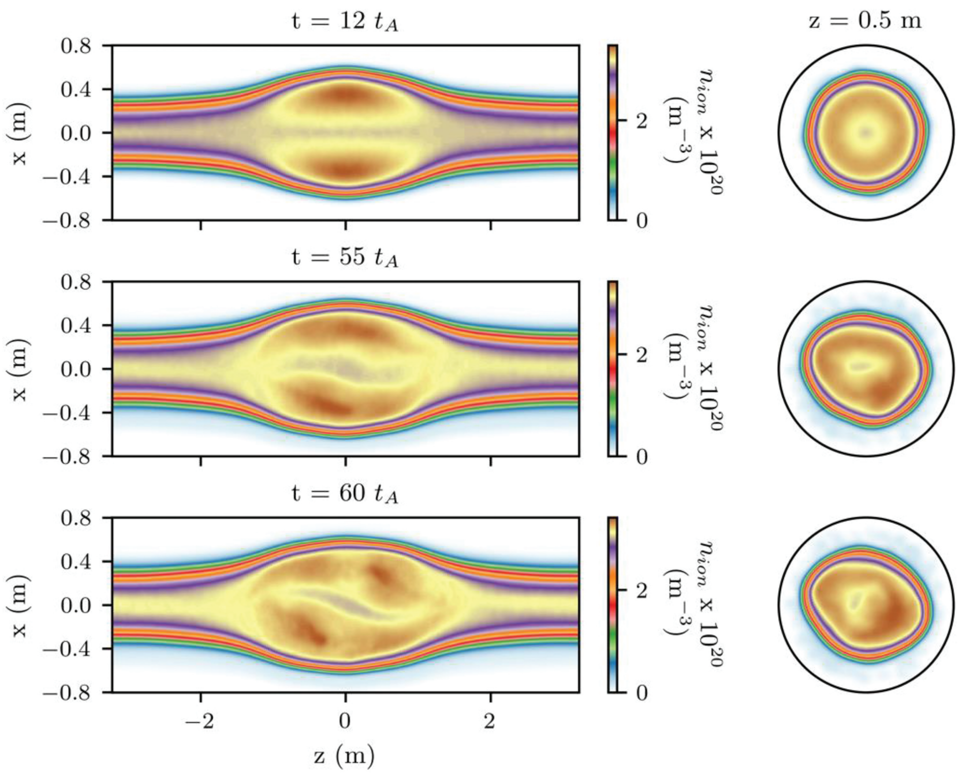

Fig. 3. Three-dimensional WarpX hybrid fluid/kinetic simulation of tilt mode in a dense FRC plasma using the new hybrid PIC solver. The initial state was an axisymmetric FRC equilibrium generated by one of TAE's multifluid equilibrium codes.[62] Evolution of the tilt mode is shown at three times: t = 12, 55, and 60 Alfvén periods after the start of the simulation.

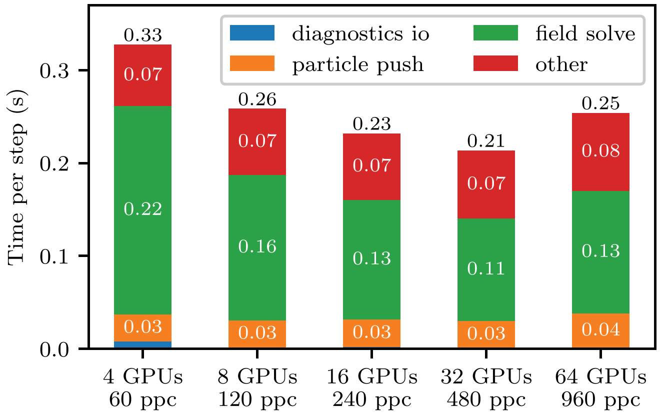

Fig. 4. Parallel scaling of the WarpX hybrid PIC algorithm on Perlmutter GPU nodes with a fixed grid size of 148 × 148 × 592, while particle count increases proportionally to the compute resources, i.e., a mixed strong and weak scaling test. Compare to Fig. 8 of Ref. [47].

Effective utilization of the OLCF Summit and the NERSC Perlmutter supercomputers has led to breakthroughs in the largest simulations using state-of-the-art GTC[65] to predict EP confinement.[66,67] Notably, EP confinement is predicted to be excellent in the ITER prefusion baseline scenario,[66] but the alpha particles in the fusion steady-state scenario suffer a large loss that requires further optimization of the operational scenario. This study is part of the SciDAC Integrated Simulation of Energetic Particles collaboration and DOE Fusion Energy Science FY2022 Theory Performance Target and was selected for two oral presentations[68,69] at the 29th Biennial International Atomic Energy Agency Fusion Energy Conference held in London in October 2023.

Leveraging its physics capabilities and computational power, GTC has been extended to simulate the FRC geometry to study turbulent transport in the FRC experiments,[24,26] a collaboration with TAE researchers that led to the development of the FRC turbulence simulation code ANC[28] and the GTC-X code.[70] In return, the new physics and computational capabilities developed through this PPP collaboration on FRC simulations benefit GTC simulations of turbulent transport and EP confinement in tokamaks and stellarators.

VI. CONCLUSION

In conclusion, the ECP has provided the platform for and has contributed to the stimulation of the frontier of fusion research to a higher plane by its vast computational expertise. And, conversely, the active fusion research stimulated the application and further growth of the supercomputing via the ECP.

Since the conclusion of the ECP, the DOE has sponsored the Post-ECP Software-Ecosystem Sustainment Project (PESO). Several ECPs, including WarpX and AMReX, have joined as founding members of a new High Performance Software Foundation. These two important developments give us the confidence to continue building our fusion plasma integrated modeling suite on top of ECP tools.

Acknowledgments

We deeply thank the co-Executive Directors of the Exascale Computing Project Industry-Agency Council, Suzy Tichenor and David Martin for their leadership and kind introduction of us into the ECP and encouragement for this research. The DOE has also provided leadership-class computing resources by way of the INCITE, ALCC, and ERCAP programs, which we have been able to efficiently utilize using ECP developed codes. Finally, the concept of PPP developed and nurtured at the DOE has provided a fruitful environment to develop this work. We also thank Rémi Lehe, Axel Huebl, and Jean-Luc Vay for their great support with WarpX. This research used resources of the NERSC, a DOE SC User Facility supported by the SC of the DOE under contract number DE-AC02-05CH11231 using NERSC award FES-ERCAP0026457.

Disclosure Statement

The authors declare no competing interests.

Funding

This work was supported by TAE Technologies, Inc., and by DOE INFUSE, with UCI contributions partially supported by DOE SciDAC.

ORCID

T. Tajima http://orcid.org/0009-0000-7850-8937
S. A. Dettrick http://orcid.org/0000-0002-1756-4625
Z. Lin http://orcid.org/0000-0003-2007-8983
R. E. Groenewald http://orcid.org/0000-0002-5282-7152
B. S. Nicks http://orcid.org/0000-0002-3876-8299
A. Veksler http://orcid.org/0009-0001-4431-4314
D. C. Barnes http://orcid.org/0000-0002-0111-1621
F. Ceccherini http://orcid.org/0000-0003-3377-6932
J. Drobny http://orcid.org/0000-0002-9733-6058
L. Galeotti http://orcid.org/0000-0003-3426-2087
S. Gupta http://orcid.org/0000-0001-5688-4122
C. K. Lau http://orcid.org/0000-0001-6702-1461
A. Necas http://orcid.org/0000-0002-9054-7653
M. Onofri http://orcid.org/0000-0003-0818-4628

References

1. D. KOTHE, S. LEE, and I. QUALTERS, “Exascale Computing in the United States,” *Comput. Sci. Eng.*, **21**, *1*, 17 (2019); http://doi.org/10.1109/MCSE.2018.2875366.
2. P. MESSINA, “The Exascale Computing Project,” *Comput. Sci. Eng.*, **19**, *3*, 63 (2017); http://doi.org/10.1109/MCSE.2017.57.
3. *Supercomputers: Algorithms, Architectures, and Scientific Computation*, F. A. MATSEN and T. TAJIMA, Eds., University of Texas Press, Austin, Texas (1986).
4. T. TAJIMA, *Computational Plasma Physics—With Applications to Fusion and Astrophysics*, Addison-Wesley, Benjamin Frontier Series, Reading, Massachusetts (1989) [Reprinted Perseus, Boulder (2004)].
5. B. I. COHEN et al., “The Numerical Tokamak Project: Simulation of Turbulent Transport,” *Comput. Phys. Commun.*, **87**, *1–2*, 1 (1995); http://doi.org/10.1016/0010-4655(94)00166-Y.
6. D. C. BARNES et al., “The Numerical Tokamak Project,” *Computing at the Leading Edge: Research in the Energy Sciences*, p. 5, A. A. MIRIN and P. T. VAN DYKE, Eds., National Energy Research Supercomputer Center (1994); http://doi.org/10.2172/10119562.
7. N. ROSTOKER et al. “Formation of a Field Reversed Configuration for Magnetic and Electrostatic Confinement of Plasma,” U.S. Patent US6891911B2 (2005).
8. S. V. PUTVINSKI, D. D. RYUTOV, and P. N. YUSHMANOV, “Fusion Reactivity of the pB11 Plasma Revisited,” *Nucl. Fusion*, **59**, *7*, 076018 (2019); http://doi.org/10.1088/1741-4326/ab1a60.
9. D. N. HILL, “Report of the FESAC Toroidal Alternates Panel,” U.S. Department of Energy (2008).
10. R. HAZELTINE, “Research Needs for Magnetic Fusion Energy Sciences.” Report of the Research Needs Workshop (ReNeW), Bethesda, Maryland, June 8-12, 2009; https://doi.org/10.2172/972502
11. R. M. MAGEE et al., “First Measurements of p11B Fusion in a Magnetically Confined Plasma,” *Nat. Commun.*, **14**, *1*, 955 (2023); http://doi.org/10.1038/s41467-023-36655-1.
12. M. W. BINDERBAUER et al. “Dynamic Formation of a Hot Field Reversed Configuration with Improved

Confinement by Supersonic Merging of Two Colliding High-Beta Compact Toroids," *Phys. Rev. Lett.*, **105**, *4*, 045003 (2010); http://doi.org/10.1103/PhysRevLett.105.045003.

13. M. TUSZEWSKI et al., "Field Reversed Configuration Confinement Enhancement Through Edge Biasing and Neutral Beam Injection," *Phys. Rev. Lett.*, **108**, *25*, 255008 (2012); http://doi.org/10.1103/PhysRevLett.108.255008.
14. H. Y. GUO et al., "Achieving a Long-Lived High-Beta Plasma State by Energetic Beam Injection," *Nat. Commun.*, **6**, 6897 (2015); http://doi.org/10.1038/ncomms7897.
15. H. GOTA et al. "Formation of Hot, Stable, Long-Lived Field-Reversed Configuration Plasmas on the C-2W Device," *Nucl. Fusion*, **59**, *11*, 112009 (2019); http://doi.org/10.1088/1741-4326/ab0be9.
16. H. GOTA et al., "Overview of C-2W: High Temperature, Steady-State Beam-Driven Field-Reversed Configuration Plasmas," *Nucl. Fusion*, **61**, *10*, 106039 (2021); http://doi.org/10.1088/1741-4326/ac2521.
17. E. A. BALTZ et al., "Achievement of Sustained Net Plasma Heating in a Fusion Experiment with the Optometrist Algorithm," *Sci. Rep.*, **7**, *1*, 6425 (2017); http://doi.org/10.1038/s41598-017-06645-7.
18. J. ROMERO, "Optimization and Feedback Control of the C-2W Field Reversed Configuration," *Proc. 29th IAEA Fusion Energy Conf.*, London, United Kingdom, October 16–21, 2023, International Atomic Energy Agency (2023).
19. J. A. ROMERO et al., "Inference of Field Reversed Configuration Topology and Dynamics During Alfvenic Transients," *Nat. Commun.*, **9**, *1*, 691 (2018); http://doi.org/10.1038/s41467-018-03110-5.
20. M. NATIONS et al. "High-Fidelity Inference of Local Impurity Profiles in C-2W Using Bayesian Tomography," *Rev. Sci. Instrum.*, **93**, *11*, 113522 (2022); http://doi.org/10.1063/5.0101741.
21. C. SCOTT et al., "Detection and Prediction of a Beam-Driven Mode in Field-Reversed Configuration Plasma with Recurrent Neural Networks," *Nucl. Fusion*, **60**, *12*, 126025 (2020); http://doi.org/10.1088/1741-4326/abb328.
22. M. DIKOVSKY et al., "Multi-Instrument Bayesian Reconstruction of Plasma Shape Evolution in the C-2W Experiment," *Phys. Plasmas*, **28**, *6*, 062503 2021); http://doi.org/10.1063/5.0049530.
23. S. A. DETTRICK et al. "Simulation of Equilibrium and Transport in Advanced FRCS," *Nucl. Fusion*, **61**, *10*, 106038 (2021); http://doi.org/10.1088/1741-4326/ac1e5f.
24. D. P. FULTON et al. "Gyrokinetic Particle Simulation of a Field Reversed Configuration," *Phys. Plasmas*, **23**, *1*, 012509 (2016); http://doi.org/10.1063/1.4930289.
25. D. P. FULTON et al. "Gyrokinetic Simulation of Driftwave Instability in Field-Reversed Configuration," *Phys. Plasmas*, **23**, *5*, 056111 (2016); http://doi.org/10.1063/1.4948285.
26. L. SCHMITZ et al., "Suppressed Ion-Scale Turbulence in a Hot High-β Plasma," *Nat. Commun.*, **7**, *1*, 13860 (2016); http://doi.org/10.1038/ncomms13860.
27. C. K. LAU et al. "Drift-Wave Stability in the Field-Reversed Configuration," *Phys. Plasmas*, **24**, *8*, 082512 (2017); http://doi.org/10.1063/1.4993630.
28. C. K. LAU et al. "Cross-Separatrix Simulations of Turbulent Transport in the Field-Reversed Configuration," *Nucl. Fusion*, **59**, *6*, 066018 (2019); http://doi.org/10.1088/1741-4326/ab1578.
29. C. LAU, F. CECCHERINI, and S. DETTRICK, "Computing Challenges in Kinetic Modeling of FRC Stability and Transport," *Bull. Am. Phys. Soc*, **64**, NM9.00007 (2019); http://meetings.aps.org/link/BAPS.2019.DPP.NM9.7.
30. C. K. LAU et al. "Electrostatic Quasi-Neutral Formulation of Global Cross-Separatrix Particle Simulation in Field-Reversed Configuration Geometry," *Phys. Plasmas*, **27**, *8*, 082504 (2020); http://doi.org/10.1063/5.0012439.
31. S. DETTRICK et al., "FPIC: A Key Next Step for Stability Studies of Advanced Beam Driven FRCs," *Bull. Am. Phys. Soc.*, **60**, BP12.00031 (2015); http://meetings.aps.org/link/BAPS.2015.DPP.BP12.31.
32. F. CECCHERINI et al., "FPIC Study of the n = 1 Toroidal Mode in FRC Plasmas," *Bull. Am. Phys. Soc.*, **61**, CP10.00087 (2016); http://meetings.aps.org/link/BAPS.2016.DPP.CP10.87
33. F. CECCHERINI, "Kinetic Investigation of Low Order Modes in FRC Plasmas" , Sherwood Fusion Theory Conf., April 4-6, 2016, Madison, Wisconsin.
34. J. T. DROBNY and D. CURRELI, "RustBCA: A High-Performance Binary-Collision-Approximation Code for Ion-Material Interactions," *J. Open Source Softw.*, **6**, *64*, 3298 (2021); http://doi.org/10.21105/joss.03298.
35. S. DETTRICK, "GPU Acceleration of a Monte Carlo Neutral Beam Injection Code Using Thrust," presented at International Conference on the Numerical Simulation of Plasmas, Long Branch, New Jersey, September 2011.
36. F. CECCHERINI et al., "Helicon and High Harmonic Fast Waves (HHFW) in FRC Configurations," *Proc. 29th IAEA Fusion Energy Conf.*, London, United Kingdom, October 16–21, 2023, International Atomic Energy Agency (2023).
37. Y. MOK, D. C. BARNES, and S. DETTRICK, "Modeling of Dynamic FRC Formation," *Bull. Am. Phys. Soc.*, **55**, GP9.00097 (2010); http://meetings.aps.org/link/BAPS.2010.DPP.GP9.97.
38. M. FOLK et al., "An Overview of the HDF5 Technology Suite and Its Applications," *Proc. EDBT/ICDT 2011 Workshop Array Databases*, Uppsala, Sweden, March 25,

2011, p. 36, Association for Computing Machinery (2011); http://doi.org/10.1145/1966895.1966900.

39. M. ONOFRI et al. "Magnetohydrodynamic Transport Characterization of a Field Reversed Configuration," *Phys. Plasmas*, **24**, *9*, 092518 (2017); http://doi.org/10.1063/1.4994681.
40. N. RATH et al. "Modeling Feedback Control of Unstable Separatrix Location in Beam-Driven Field-Reversed Configurations," *Phys. Plasmas*, **24**, *4*, 042504 (2017); http://doi.org/10.1063/1.4979188.
41. M. ONOFRI et al., "Transport Simulations of the C-2 and C-2U Field Reversed Configurations with the Q2D Code," *Bull. Am. Phys. Soc.*, **61**, CP10.00088 (2016); http://meetings.aps.org/link/BAPS.2016.DPP.CP10.88.
42. M. C. THOMPSON et al. "Integrated Diagnostic and Data Analysis System of the C-2W Advanced Beam-Driven Field-Reversed Configuration Plasma Experiment," *Rev. Sci. Instrum.*, **89**, *10*, 10K114 (2018); http://doi.org/10.1063/1.5037693.
43. J. A. STILLERMAN et al. "MDSplus Data Acquisition System," *Rev. Sci. Instrum.*, **68**, *1*, 939 (1997); http://doi.org/10.1063/1.1147719.
44. T. FREDIAN et al. "MDSplus Yesterday, Today and Tomorrow," *Fusion Eng. Des.*, **127**, 106 (2018); http://doi.org/10.1016/j.fusengdes.2017.12.010.
45. J. AHRENS, B. GEVECI, and C. LAW, *ParaView: An End-User Tool for Large Data Visualization*, Elsevier (2005).
46. H. CHILDS et al., "VisIt: An End-User Tool for Visualizing and Analyzing Very Large Data," *High Performance Visualization–Enabling Extreme-Scale Scientific Insight*, p. 357 (2012).
47. R. E. GROENEWALD et al. "Accelerated Kinetic Model for Global Macro Stability Studies of High-Beta Fusion Reactors," *Phys. Plasmas*, **30**, *12*, 122508 (2023); http://doi.org/10.1063/5.0178288.
48. S. ABHYANKAR et al., "PETSc/TS: A Modern Scalable ODE/DAE Solver Library," arXiv:1806.01437 (2018); http://doi.org/10.48550/arXiv.1806.01437.
49. W. F. GODOY et al. "ADIOS 2: The Adaptable Input Output System. A Framework for High-Performance Data Management," *SoftwareX*, **12**, 100561 (2020); http://doi.org/10.1016/j.softx.2020.100561.
50. W. ZHANG et al., "AMReX: Block-Structured Adaptive Mesh Refinement for Multiphysics Applications," *Int. J. High Perform. Comput. Appl.*, **35**, *6*, 508 (2021); http://doi.org/10.1177/10943420211022811.
51. ECP-WarpX/WarpX; https://github.com/ECP-WarpX/WarpX (current as of June 11, 2024).
52. D. C. BARNES, "Time-Explicit Darwin PIC Algorithm," *J. Comput. Phys.*, **462**, 111151 (2022); http://doi.org/10.1016/j.jcp.2022.111151.
53. T. TAJIMA and J. M. DAWSON, "Laser Electron Accelerator," *Phys. Rev. Lett.*, **43**, *4*, 267 (1979); http://doi.org/10.1103/PhysRevLett.43.267.
54. T. TAJIMA, K. MIMA, and J. M. DAWSON, "Alfvén Ion-Cyclotron Instability: Its Physical Mechanism and Observation in Computer Simulation," *Phys. Rev. Lett.*, **39**, *4*, 201 (1977); http://doi.org/10.1103/PhysRevLett.39.201.
55. T. TAJIMA and J. M. DAWSON, "Ion Cyclotron Resonance Heating and the Alfvén-Ion Cyclotron Instability," *Nucl. Fusion*, **20**, *9*, 1129 (1980); http://doi.org/10.1088/0029-5515/20/9/012.
56. B. S. NICKS, S. PUTVINSKI, and T. TAJIMA, "Stabilization of the Alfvén-Ion Cyclotron Instability Through Short Plasmas: Fully Kinetic Simulations in a High-Beta Regime," *Phys. Plasmas*, **30**, *10*, 102108 (2023); http://doi.org/10.1063/5.0163889.
57. D. W. SWIFT, "Use of a Hybrid Code for Global-Scale Plasma Simulation," *J. Comput. Phys.*, **126**, *1*, 109 (1996); http://doi.org/10.1006/jcph.1996.0124.
58. E. BELOVA et al., "Numerical Study of Global Stability of Oblate Field-Reversed Configurations," *Phys. Plasmas*, **8**, *4*, 1267 (2001); http://doi.org/10.1063/1.1355027.
59. E. V. BELOVA et al., "Kinetic Effects on the Stability Properties of Field-Reversed Configurations. I. Linear Stability," *Phys. Plasmas*, **10**, *6*, 2361 (2003); http://doi.org/10.1063/1.1568118.
60. E. V. BELOVA et al., "Kinetic Effects on the Stability Properties of Field-Reversed Configurations. I. Nonlinear Evolution," *Phys. Plasmas*, **11**, *5*, 2523 (2004); http://doi.org/10.1063/1.1666293.
61. A. LE et al. "Hybrid Simulations of Magnetic Reconnection with Kinetic Ions and Fluid Electron Pressure Anisotropy," *Phys. Plasmas*, **23**, *3*, 032114 (2016); http://doi.org/10.1063/1.4943893.
62. L. GALEOTTI et al., "Plasma Equilibria with Multiple Ion Species: Equations and Algorithm," *Phys. Plasmas*, **18**, *8*, 082509 (2011); http://doi.org/10.1063/1.3625275.
63. R. M. HEDLOF et al., "Physical Effects of Aleph's Semi-Implicit Particle-in-Cell Scheme," *Proc. IEEE Int. Conf. Plasma Science (ICOPS)*, Santa Fe, New Mexico, May 21–25, 2023, IEEE (2023); http://doi.org/10.1109/ICOPS45740.2023.10481040.
64. A. NECAS et al. "Simulation of Coherent Terahertz Radiation from Relativistic Laser-Solid Interactions," *Proc. 63rd Annual Mtg. APS Division of Plasma Physics*, Pittsburgh, Pennsylvania, November 8–12, 2021, American Physical Society (2021).
65. Z. LIN et al. "Turbulent Transport Reduction by Zonal Flows: Massively Parallel Simulations," *Science*, **281**, *5384*, 1835 (1998); http://doi.org/10.1126/science.281.5384.1835.

66. G. BROCHARD et al., “Saturation of Fishbone Instability by Self-Generated Zonal Flows in Tokamak Plasmas,” *Phys. Rev. Lett.*, **132**, *7*, 075101 (2024); http://doi.org/10.1103/PhysRevLett.132.075101.

67. P. LIU et al., “Regulation of Alfvén Eigenmodes by Microturbulence in Fusion Plasmas,” *Phys. Rev. Lett.*, **128**, *18*, 185001 (2022); http://doi.org/10.1103/PhysRevLett.128.185001.

68. G. BROCHARD et al., “Saturation of Fishbone Instability by Self-Generated Zonal Flows in Tokamak Plasmas,” *Proc. 29th IAEA Fusion Energy Conf.*, London, United Kingdom, October 16–21, 2023, International Atomic Energy Agency (2023).

69. Z. LIN et al., “Prediction of Energetic Particle Confinement in ITER Operation Scenarios,” *Proc. 29th IAEA Fusion Energy Conf.*, London, United Kingdom, October 16–21, 2023, International Atomic Energy Agency (2023).

70. W. H. WANG et al. “Effects of Equilibrium Radial Electric Field on Ion Temperature Gradient Instability in the Scrape-Off Layer of a Field-Reversed Configuration,” *Plasma Phys. Controlled Fusion*, **63**, *6*, 065001 (2021); http://doi.org/10.1088/1361-6587/abf403.