Integrated core-pedestal modeling with the AToM framework

J. Candy[1], E. Belli[1], D. Green[3], C. Holland[2], O. Meneghini[1], J.M. Park[3], S. Smith[1]

and the AToM team.

[1]*General Atomics, San Diego, California, USA*

[2]*University of California at San Diego, San Diego, CA*

[3]*Oak Ridge National Laboratory, Oak Ridge, TN*

Introduction The aim of this work is to report on integrated simulation capabilities of the *Advanced Tokamak Modeling* (AToM) framework[1], as well as the leadership computing effort that provides the theoretical physics basis for the modeling capability. The original AToM[0] SciDAC-3 project was a 3-year pilot project that concluded in 2017. For the period 2017-2022, a new 5-year AToM SciDAC-4 project has begun. The overall scope of AToM is broad, with six distinct research thrusts; namely, (1) software environment, performance and packaging, (2) physics component integration, (3) validation and uncertainty quantification, (4) physics and scenario exploration, (5) data and metadata management, and (6) liaisons to other SciDAC partnerships. For thrust (2), we support two separate but comprehensive approaches via the OMFIT-TGYRO [1] and the IPS-FASTRAN [2] workflows based on the OMFIT[2] and IPS frameworks. Both workflows account for key physics processes including the strong interplay between core turbulent-plus-collisional transport (determined by the core profiles), pedestal structure, current profile and plasma equilibrium. Both workflows calculate the self-consistent, steady-state solution to this strongly-coupled problem, and thereby provide a state-of-the-art capability for predictive tokamak modeling. In both workflows, the quality of the physics prediction is limited by the accuracy of the core transport fluxes as determined by TGLF [3] (turbulent flux) and NEO [4, 5] (neoclassical flux, and bootstrap current), and the pedsetal width and height as determined by EPED [6]. Importantly, TGLF is itself *calibrated* to the nonlinear fluxes obtained by GYRO code, and future plans include comprehensive recalibration based on multiscale simulations from the CGYRO [7] code.

The OMFIT-TGYRO workflow For the OMFIT-TGYRO workflow, which makes heavy use of the GACODE[3] suite of codes, kinetic plasma equilibrium is computed by EFIT using profile data from the TGYRO transport solver and bootstrap current from direct kinetic solutions provided by the NEO neoclassical module. Given an EFIT equilibrium, core profiles are evolved

[1]https://scidac.github.io/atom/
[2]https://gafusion.github.io/OMFIT-source/
[3]https://gafusion.github.io/doc/

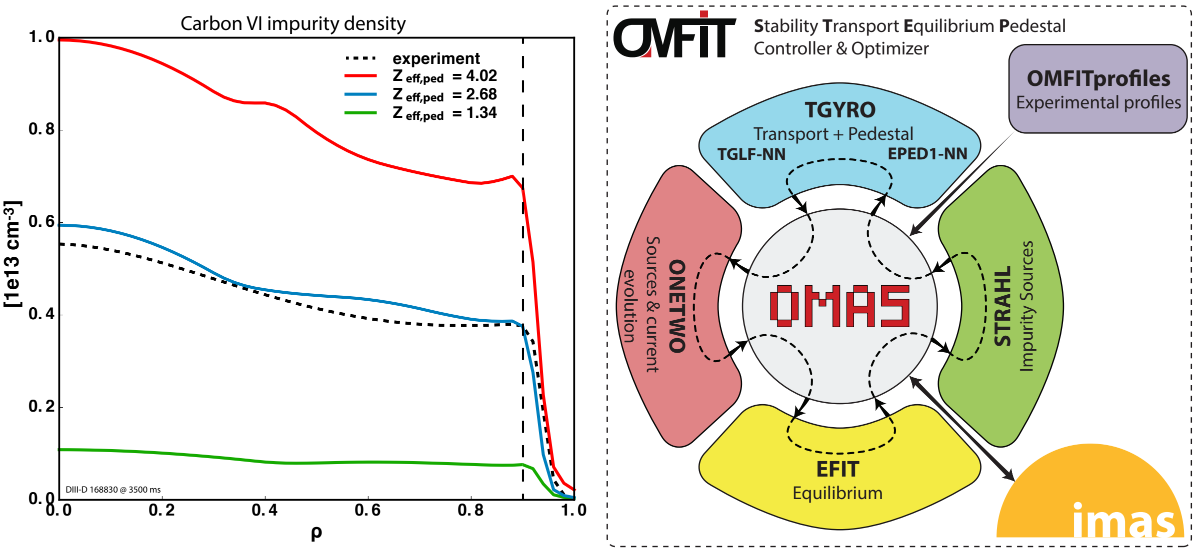

Figure 1: Impurity profiles generated by OMFIT-TGYRO-STRAHL coupling (left plot). Conceptual illustration of function of the role of OMAS in this new integrated workflow (right plot).

by the transport solver by combining collisional (NEO) and turbulent (TGLF) fluxes to maintain balance with heating and fueling sources (ONETWO). The updated value of the global plasma pressure is passed to EPED to obtain the self-consistent pedestal structure. EPED combines criticality constraints for larger-scale peeling-ballooning and smaller-scale kinetic ballooning modes to determine pedestal height and width. The scheme is iterated to convergence in a few iterations and is independent of initial profiles. By using only the electron density at the top of the pedestal as an input, profiles are calculated from magnetic axis to separatrix, in good agreement with experiment [1].

Reduced models via neural-network approximation Significant effort within AToM has been devoted to development of reduced models based on machine learning techniques. Specifically, neural-networks have been used to generate an accurate functional representation of models for pedestal structure (EPED1-NN) and turbulent transport physics (TGLF-NN). These models have been tightly coupled within the TGYRO transport solver. Initial application of the neural-network approach to realistic DIII-D core-pedestal coupled simulations has proven to be robust and remarkably efficient computationally [1]. Recent work in this area aims to extend the original TGLF-NN model to handle multiple ion species, as necessary for modeling impurity transport (deuterium and carbon fluxes to for DIII-D plasmas, and a deuterium-tritium mix plus helium ash for ITER plasmas). To evolve the impurity density, TGYRO is coupled within OMFIT to the 1D impurity transport code STRAHL, as illustrated in Fig. 1. Here, TGYRO provides the transport fluxes that are used to calculate the diffusion and pinch profiles used by STRAHL. The resulting workflow evolves non-trace impurities and fully accounts for the feedback between impurity dilution and gradients/fluxes of the main ions. Notably, the workflow is fully compatible with the ITER Integrated Modeling and Analysis Suite (IMAS). This technical achievement was enabled by transferring information betweem physics components within OMFIT by means of the OMAS library. At the heart of OMAS is the idea of providing a convenient API that can store data in a format compatible with the IMAS data model, but using other storage systems in addition to the one provided by IMAS itself.

CGYRO: providing the basis for TGLF The quality of profile prediction for an integrated modeling workflow is only as good as the quality of the reduced model of core turbulent flux. At the heart of all AToM profile-prediction capabilities is TGLF [3], which uses a linear gyro-Landau-fluid eigenvalue solver coupled with a sophisticated saturation rule. The saturation rule is derived exclusively from an ensemble (i.e., a database) of nonlinear GYRO/CGYRO gyrokinetic simulation. Over the last decade, this database has been populated by GYRO simulations for core plasma parameters. However, this database resolves only long-wavelength turbulence for which $k_\perp \rho_i < 1$. To properly capture dynamics of ETG turbulence and potential multiscale corrections, future simulations used for TGLF calibration will be carried out with CGYRO, which has been highly optimized (both algorithmically and computationally) for multiscale simulations. A significant amount of recent optimization work has been carried out for both multicore (Xeon Phi) and GPU (Kepler/Volta) systems. Representative scaling results are shown in Fig. 2. An example of a CGYRO multiscale simulation of a DIII-D ITER baseline scenario discharge is given in Fig. 3, illustrating the necessity of multiscale simulation (versus traditional ion-scale simulation) for the outer radii of ITER-type plasmas.

Work supported by the U.S. Department of Energy under DE-SC0017992, DE-FC02-06ER54873, DE-FG02-95ER54309 and DE-FC02-04ER54698.

[1] O. Meneghini, S. Smith, *et al.*, *Nuclear Fusion*, **57** (2017) 0866034
[2] J. Park, J. Ferron, C. Holcomb, *et al.*, *Phys. Plasmas*, **25** (2018) 012506
[3] G. Staebler, J. Kinsey, and R. Waltz, *Phys. Plasmas*, **12** (2005) 102508
[4] E. Belli and J. Candy, *Plasma Phys. Control. Fusion*, **50** (2008) 095010
[5] E. Belli and J. Candy, *Plasma Phys. Control. Fusion*, **54** (2012) 015015
[6] P. Snyder, *et al.*, *Nucl. Fusion*, **51** (2011) 103016
[7] J. Candy, E. Belli, and R. Bravenec, *J. Comput. Phys.*, **324** (2016) 73

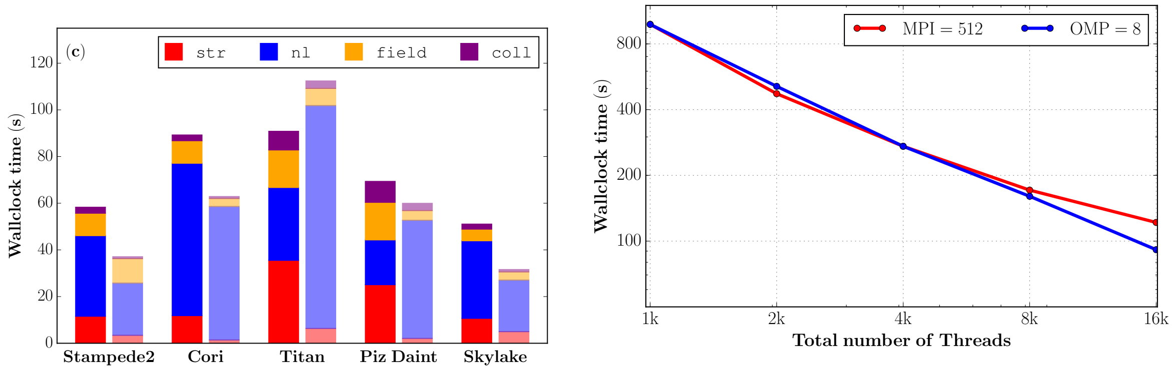

Figure 2: CGYRO performance data. Left plot: kernel performance plots for each of the four CGYRO kernels (streaming (str), nonlinear (nl) Maxwell equations (field) and collisions (coll) on five leadership computers. The solid left bars for each system denote computation times, whereas the faded right bars denote communication time. Right plot: hybrid MPI scaling for the `nl03` test case on Cori KNL. The red curve show scaling at fixed MPI task count (from 2 OMP threads on 8 nodes to 32 OMP threads on 128 nodes). The blue curve shows scaling at fixed OMP thread count (from 128 MPI tasks at 8 nodes to 2048 MPI tasks on 128 nodes). Thus there is a nearly a perfect tradeoff between OMP and MPI, except for the last OMP (red) scaling point (32 OMP threads).

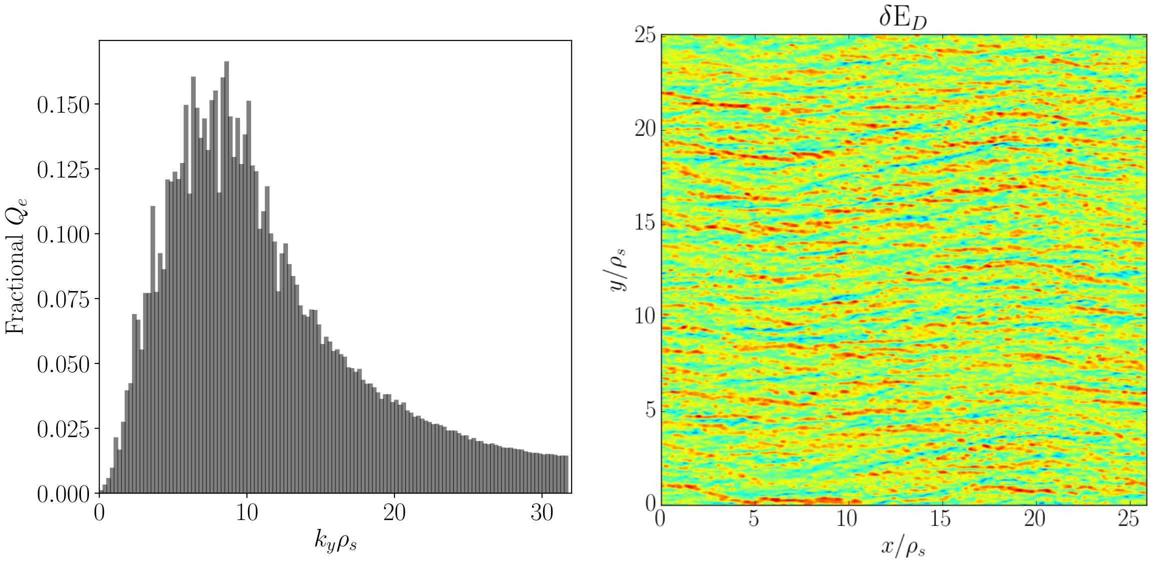

Figure 3: CGYRO simulation of a DIII-D ITER baseline scenario discharge (shot 164988) at $r/a = 0.92$. Left plot: Fractional contribution to electron energy flux, Q_e/Q_{GB}, as a function of binormal wavenumber k_y. Total electron flux is equal to the area under the curve. Right plot: snapshot of deuterium energy fluctuations. Simulation results in this case are in close agreement with experiment and illustrate necessity of multiscale resolution to describe ITER-relevant scenarios.