

Integrated Fusion Simulation with Self-Consistent Core-Pedestal Coupling

O. Meneghini,* P.B. Snyder, S.P. Smith, J. Candy, G.M. Staebler, E.A. Belli, and L.L. Lao

General Atomics, San Diego CA

J.M. Park, D.L. Green, and W. Elwasif

Oak Ridge National Laboratory, Oak Ridge TN

B.A. Grierson

Princeton Plasma Physics Laboratory, Princeton NJ

C. Holland

University of California San Diego, San Diego CA

(Dated: March 1, 2016)

Abstract

Accurate prediction of fusion performance in present and future tokamaks requires taking into account the strong interplay between core transport, pedestal structure, current profile and plasma equilibrium. An integrated modeling workflow capable of calculating the steady-state self-consistent solution to this strongly-coupled problem has been developed. The workflow leverages state-of-the-art components for collisional and turbulent core transport, equilibrium and pedestal stability. Validation against DIII-D discharges shows that the workflow is capable of robustly predicting the kinetic profiles (electron and ion temperature and electron density) from the axis to the separatrix in good agreement with the experiments. An example application is presented, showing self-consistent optimization for the fusion performance of the 15 MA D-T ITER baseline scenario as functions of the pedestal density and ion effective charge Z_{eff} .

* meneghini@fusion.gat.com

I. INTRODUCTION

The physical processes that govern the core and pedestal regions of tokamak plasmas are strongly coupled. The peeling-balloonning (PB) and kinetic ballooning modes (KBM) impose constraints on the pedestal structure, which in turn strongly affect core confinement (or equivalently, core temperature and pressure) in H-mode plasmas. Because PB stability is impacted by the global Shafranov shift, which is proportional to the core pressure, this leads to a feedback cycle between core and pedestal [1, 2]. Similarly, the strength and scaling of collisional (neoclassical) and turbulent core transport depend not only on the temperature and density profiles, but also on the plasma equilibrium and current profiles. Steady-state core profiles are achieved when the transport processes balance the heating and particle sources. Thus, predictive modeling of profiles from core to edge represents successful coupling of (a) plasma equilibrium, (b) core turbulence and transport, (c) particle and energy sources, and (d) pedestal structure.

Although the interaction between the H-mode pedestal and the core plasma has long been recognized as an important effect [1–3], it has not generally been accounted for in a self-consistent way. Typical transport studies keep the pedestal shape fixed while equilibrium, core transport and sources are iterated self-consistently. This iterative process is illustrated in Fig. 1a. In this traditional approach, the pedestal pressure and width is most commonly assumed to be known and fixed [4–7]. Other approaches are to calculate the pedestal parameters from scaling laws [3, 8–10] or to employ a pedestal structure model [11, 12] such as EPED [2, 13, 14] *without* self-consistency. In the latter case, the stabilizing effect of the equilibrium Shafranov shift caused by the core pressure is assumed *ab initio*. Although this approach may be acceptable for transport analysis of existing experiments, for which the pedestal boundary condition is known, it can lead to large uncertainties for true predictive modeling of future devices [9]. The lack of self-consistency means that the predicted total plasma pressure (labeled as $\beta_{n,out}$ in Fig. 1) changes based on the total plasma pressure that is input to the pedestal model ($\beta_{n,in}$). In other words, the outcome of the core transport simulation depends on the initial assumptions of the pedestal or the inputs to the pedestal structure model.

In this paper, we will refer to two different flux-surface labels. The first is a geometric quantity, r , which is the half-width of the flux surface measured at the elevation of the

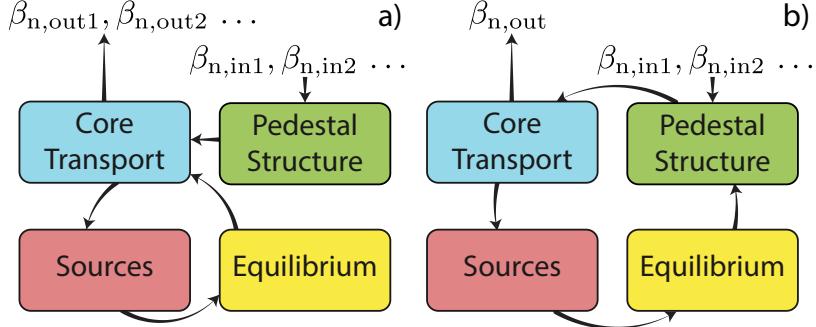


FIG. 1. High-level schematics comparing (a) the traditional workflow used for predictive simulations in which the pedestal structure is held fixed (and the outcome of the simulation depends on the initial assumptions made in the pedestal structure calculation), and (b) the workflow presented in the present paper for which the pedestal structure calculation is part of the iteration loop (and the simulation outcome is independent of the initial assumptions made in the initial pedestal structure calculation).

centroid [5]. We remark that r is sometimes referred to as the *Miller minor radius* and is the fundamental radial variable in numerous modeling codes. This definition is suitable for a plasma cross-section of arbitrary shape and elevation, and in the limit of an unshifted circular plasma, r reduces to the radius of the circle. In general it is computed numerically after the plasma equilibrium has been determined. The second flux-surface label is the *square root of the normalized toroidal flux*,

$$\rho \doteq \sqrt{\frac{\chi_t(r)}{\chi_t(a)}}, \quad (1)$$

where a is the value of r at the last closed flux surface (LCFS). The quantity ρ is commonly used for plasma modeling, but it should be noted that it does *not* correspond to a physical length. Throughout this paper we refer to the volume-averaged normalized plasma pressure $\beta_n \doteq \beta_t/I_n$ [15] where

$$\beta_t = \frac{\langle p \rangle}{B_t^2/2\mu_0} \quad \text{and} \quad I_n = \frac{I_p}{aB_t}. \quad (2)$$

Here, $\langle p \rangle$ is the volume-averaged plasma pressure in Pa, B_t is the vacuum magnetic field strength in T, I_p is the plasma current expressed in MA, and a is the radius in m. In this paper, we will always express β_n as a percentage.

As part of the SciDAC AToM project [16] we have developed an iterative workflow to calculate the steady-state self-consistent solution to the coupled core-pedestal problem. This

new workflow, as illustrated in Fig. 1b, implements an iteration loop that couples (a) plasma equilibrium, (b) core turbulence and transport, (c) particle and energy sources, and (d) pedestal structure. The key innovation is to update the global pressure (derived from the core-transport prediction at the previous step) that is input to the pedestal model. We remark that for the cases presented in this paper, that the converged solution to the coupled system is found to be unique; that is, independent of the initial guess of β_n .

In Sec. II we will describe in greater detail the physics modules implemented in the new workflow, as well as the coupling strategies used to connect them. The implementation of the iterative workflow, which is used to find the steady-state self-consistent solution, is given in Sec. IV. Further, in Sec. IV, we illustrate the workflow by simulating a DIII-D ITER baseline scenario discharge. The insensitivity of the self-consistent solution to different initial guesses for the global pressure is also clearly demonstrated. Finally, in Sec. V, we further apply the iterative workflow to perform a self-consistent optimization of the fusion performance for the 15 MA D-T ITER baseline scenario as function of the pedestal density $n_{e,ped}$ and ion effective charge $Z_{eff,ped}$.

II. CORE TRANSPORT CALCULATION METHODOLOGY

In the present work, profiles are determined by matching transport fluxes with heating and particle sources using the TGYRO [5] transport module. In addition to its capability as a transport code, TGYRO is also a parallel transport manager with the ability to call multiple concurrent instances of the TGLF [17] transport model, or more generally, massively-parallel GYRO [18] simulations. TGYRO can simultaneously manage the parallel execution of kinetic NEO [19–21] simulations. The total transport fluxes are given by a sum of the NEO neoclassical fluxes and the turbulent fluxes (via TGLF in this paper). The TGYRO approach (combining TGLF and NEO) has been previously applied with great success for the prediction of kinetic profiles in the core region (see Fig. 2) of tokamak plasmas [4, 5, 22–26].

A. Neoclassical flux and bootstrap current: NEO

Although formulae exist for the neoclassical fluxes and the bootstrap current, the direct kinetic approach used by NEO ensures the highest possible accuracy for the neoclassical processes. NEO solves the drift-kinetic equation,

$$v_{\parallel} \mathbf{b} \cdot \nabla g_{1a} - \sum_b C_{ab}^L(g_{1a}, g_{1b}) = -\mathbf{v}_D \cdot \nabla f_{0a} - \frac{z_a e}{T_{0a}} f_{0a} \vec{v}_D \cdot \nabla \Phi_0 , \quad (3)$$

where g_{1a} is the nonadiabatic perturbed distribution function. In NEO, no approximations beyond the drift-ordering are made. Full sonic toroidal rotation and centrifugal terms are retained, and general flux-surface shape is treated. The full linearized Fokker-Planck collision operator is used for the collision dynamics, with complete cross-species collisional coupling for arbitrary mass ratio and an arbitrary number of ion species, thus strictly preserving ambipolarity. NEO has been extensively benchmarked with analytic theories, as well as with NCLASS [27], over a wide range of parameters and in various asymptotic limits [19, 20, 28]. With NEO, we maintain accuracy even in the case of strong plasma shaping, large trapped fraction, multiple ions species, sonic rotation, fast particles, and high pedestal collisionality. We note however that the cost of this approach scales with the square of the number of species because of the complex interspecies collisional coupling. From the solution of the distribution function, we can compute the fluxes and bootstrap current via

$$\Gamma_a = \left\langle \int d^3v g_{1a} \mathbf{v}_D \cdot \nabla r \right\rangle , \quad (4)$$

$$Q_a = \left\langle \int d^3v m_a \varepsilon g_{1a} \mathbf{v}_D \cdot \nabla r \right\rangle , \quad (5)$$

$$\langle j_{\parallel} B \rangle = \sum_a z_a e \left\langle B \int d^3v g_{1a} v_{\parallel} \right\rangle . \quad (6)$$

B. Turbulent flux: TGLF and GYRO

Rather than attempting direct gyrokinetic simulation to determine the particle and energy transport coefficients, we can achieve a speedup of millions by instead using the Trapped Gyro Landau Fluid (TGLF) transport model [4]. This is a quasilinear transport model using the same general methods as the Weiland [29] and GLF23 [30] models. The linear eigenvalues and eigenfunctions of a system of trapped gyro-Landau fluid equations [31] are used to evaluate the quasilinear weights of the transport fluxes. The linear growth rate

spectrum is used to compute a model amplitude for the saturated turbulence. TGLF extends its predecessor GLF23 by having a more accurate system of equations valid at both electron and ion scales. TGLF bridges the gap between electron and ion scale instabilities and has a more accurate trapped particle treatment. The linear eigenvalues of TGLF have been extensively benchmarked with gyrokinetic calculations [31]. The model for the saturation of the turbulence has been fit to a database of nonlinear gyrokinetic turbulence simulations with the GYRO code. It is important to emphasize that no fitting parameters have been adjusted to experimental data. Instead, TGLF is designed for approximate validation of gyrokinetic turbulent transport with larger datasets than is possible with full gyrokinetic simulation. The TGLF equations are fully electromagnetic but are limited to high toroidal mode number instabilities due to the gyrokinetic ordering assumptions. Only electron pitch angle scattering collisions are included in TGLF at present. A new saturation model is being developed for TGLF to more accurately represent the electron scale saturation of the turbulence and its coupling to ion scales (Staebler 2016, in press).

C. Steady-state profile prediction: TGYRO

Evolution of the plasma profiles (density, temperature, electric field) occurs slowly, on the transport timescale, in response to the second-order collisional (NEO) and turbulent (TGLF) radial transport of particles, energy and momentum. First, the equation for **density evolution** solved in TGYRO is

$$\frac{\partial \langle n_a \rangle}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial r} (V' \Gamma_a) = S_{n,a} , \quad (7)$$

where V is the volume enclosed by the flux surface, $V' = \partial V / \partial r$, and r is the Miller minor radius defined in Sec. I. The sources and fluxes, defined in Table I, are

$$S_{n,a} = S_{n,a}^{\text{beam}} + S_{n,a}^{\text{wall}} , \quad (8)$$

$$\Gamma_a = \Gamma_a^{\text{neo}} + \Gamma_a^{\text{tur}} . \quad (9)$$

Typically, only the electron density n_e is evolved and the ions are updated according to quasineutrality. This is a practical approach to overcome the limitation imposed by uncertainties in the ion particle sources. Next, the equation for the **pressure evolution** (which

is effectively the equation for temperature evolution) of species a is

$$\frac{\partial \langle W_a \rangle}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial r} (V' Q_a) + \Pi_a \frac{\partial \omega_0}{\partial \psi} = S_{W,a} , \quad (10)$$

where ω_0 is the plasma rotation, and the sources and fluxes are

$$S_{W,a} = S_{W,a}^{\text{aux}} + S_{W,a}^{\text{rad}} + S_{W,a}^{\alpha} + S_{W,a}^{\text{tur}} + S_{W,a}^{\text{col}} , \quad (11)$$

$$Q_a = Q_a^{\text{neo}} + Q_a^{\text{tur}} . \quad (12)$$

Finally, the equation for **momentum evolution** (which is effectively the equation for the radial electric field evolution) is

$$\frac{\partial}{\partial t} \left(\omega_0 \langle R^2 \rangle \sum_a m_a n_a \right) + \frac{1}{V'} \frac{\partial}{\partial r} \left(V' \sum_a \Pi_a \right) = \sum_a S_{\omega,a} , \quad (13)$$

where the total momentum flux is

$$\Pi_a = \Pi_a^{\text{neo}} + \Pi_a^{\text{tur}} . \quad (14)$$

In the pressure and momentum equations, we have introduced the *toroidal rotation frequency* $\omega_0 = cE_r/(RB_p)$ [20] where E_r is the radial electric field and B_p is the poloidal magnetic field. This approach is valid in the limit that the rotation is dominantly toroidal and *sonic*, which typically requires significant beam heating power to achieve. When the rotation is diamagnetic – that is, in the spontaneous rotation regime – an accurate calculation of the electric field requires a higher-order treatment [32] and is therefore beyond the capabilities of the present modeling approach.

The numerical approach used in TGYRO is motivated by an essential feature of profile evolution. In steady-state, the time derivatives in the transport equations are exactly zero, and the balance of transport fluxes (particle, energy and momentum) with heating sources is a complicated nonlinear root finding problem. Obviously, for time-independent heating sources, there is a unique set of profiles for which steady-state is achieved. Importantly, even when the plasma is not in steady state, the physical rate of evolution of profiles typically remains slow, so that the time-evolution should be viewed as almost quasi-static. In the present paper we use TGYRO to compute, at every iteration level, the flux-matching profiles, in which case the time derivatives in the TGYRO transport equations are ignored. The case of time-dependent sources, though, is straightforward and can be retained in the general case.

In the more general case, the steady-state solver becomes the kernel for a fully-implicit time advance.

As described in Ref. [5], TGYRO uses a novel method to solve the problem. The transport equations are integrated over volume to yield an algebraic equation expressing the balance of fluxes and sources. For example, the energy equations becomes

$$Q_a(r) = \frac{1}{V'(r)} \int_0^r dx V'(x) S_{W,a}(x) \doteq Q_a^T(r) , \quad (15)$$

Variable	Definition
Γ_a^{tur}	Turbulent particle flux
Γ_a^{neo}	Neoclassical particle flux
Q_a^{tur}	Turbulent energy flux
Q_a^{neo}	Neoclassical energy flux
Π_a^{tur}	Turbulent momentum flux
Π_a^{neo}	Neoclassical momentum flux
S_a^{aux}	Auxiliary heating power density
S_a^{rad}	Radiation (loss) power density
S_a^α	Alpha-particle heating power density
S_a^{tur}	Turbulent exchange power density
S_a^{col}	Collisional exchange power density

TABLE I. Definitions of the sources and fluxes appearing the TGYRO transport equations.

We refer to the quantity Q_a^T as the *target flux*, since that is the flux we will try to match by adjusting parameters in the TGLF and NEO calculations. Since the latter codes require the profile gradients as inputs, we define the logarithmic gradients $z_a \doteq -(1/T_a)\partial T_a/\partial r$. Then, if we specify the temperature at an arbitrary *matching radius* r_* (normally chosen in the vicinity of the pedestal) as $T_{a*} \doteq T_a(r_*)$, the gradients uniquely determine the *smooth* temperature profiles, T_a :

$$T_a(r) = T_{a*} \exp \left(\int_r^{r_*} dx z_a(x) \right) , \quad (16)$$

where z_a is taken to be piecewise-linear. On a discrete grid $\{r_j\}$, the smooth temperature profile coincides with the discrete values obtained by the trapezoidal rule

$$T_a(r_{j-1}) = T_a(r_j) \exp \left\{ \left[\frac{z_a(r_j) + z_a(r_{j-1})}{2} \right] [r_j - r_{j-1}] \right\}. \quad (17)$$

To put the problem into discrete form, we define a vector of gradients (independent variables) $z_{a,j} = z_a(r_j)$, transport fluxes, $Q_{a,j} = Q_a(r_j)$ and target fluxes $Q_{a,j}^T = Q_a^T(r_j)$. Then, the equations to be solved are

$$Q_{a,j} = Q_{a,j}^T, \quad (18)$$

Note that $Q_{a,j}$ is computationally very expensive to evaluate, whereas $Q_{a,j}^T$ are very fast and take an insignificant amount of compute time. The flux-matching calculations are performed at a few (typically less than 10) radial locations in the plasma, using a sparse Newton method approach.

We remark that TGYRO also operates on an arbitrary (irregular) radial grid, and can function with only a single flux-matching point $r_1 > 0$ where the transport fluxes are matched to the sources. TGYRO further uses a novel method to treat the magnetic axis which allows the simulation to avoid the breakdown of the drift-ordering which occurs at $r \rightarrow 0$. This breakdown is illustrated in Fig. 15 of Ref. [19], and is physically connected to the transformation of banana orbits into potato orbits. In other words, the underlying assumption that the orbit width is small compared to the gradient scale length is not valid as $\rho/a \rightarrow 0$. But, because the gradient scale-length must remain large and well-behaved, in TGYRO we simply assume a linear profile of gradient scale-length from the magnetic axis to the first flux-matching point. Then, the location of the first flux-matching point can be taken just outside the thermal potato radius. This flexibility can also be used to facilitate simulation of discharges which are affected by sawtooth instabilities. In these cases, the $q \simeq 1$ region of the plasma is dominated by the rapid dynamics of MHD-driven sawtooth instabilities. But again, the gradients profiles ought to be smooth as the magnetic axis is approached, so that first flux-matching point can be placed just outside the sawtooth region.

III. PEDESTAL STABILITY AND STRUCTURE

In the pedestal region, the very short radial scales associated with the equilibrium present a substantial challenge to traditional theoretical approaches. Because the observed fluctu-

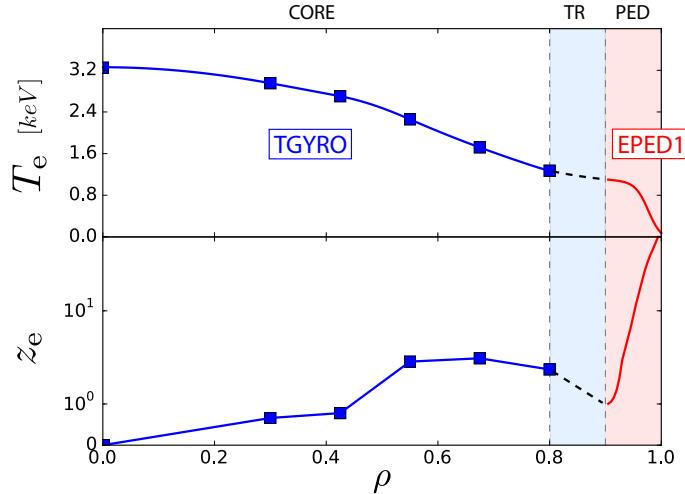


FIG. 2. Illustration of radial connectivity of core (CORE), transition (TR) and pedestal (PED) regions during self-consistent modeling, as described in detail in Sec. II. In the CORE region, TGYRO computes the gradient scale-length profile (bottom frame) $z_e = -d \ln T_e / dr$ using non-linear root-finding, and obtains the smooth temperature profile (top frame) by integration. The introduction of a finite-sized transition region (TR) allows smooth matching of core profiles to pedestal profiles for cases in which the core model may underestimate the transport.

ation scales overlap strongly with the equilibrium scale, strictly speaking it is not possible to rigorously employ the approach of traditional transport theory. The assumptions of the traditional theory that are violated are 1) that the fluctuation scales are analytically separated from the equilibrium and transport scales, 2) that the fluctuation scale is treated with separate gyrokinetic and neoclassical codes, and 3) that transport and equilibrium scale physics can be evaluated with transport solvers and magnetohydrodynamic (MHD) codes. Thus, a different approach is required.

A. The EPED1 pedestal model

Despite the theoretical challenge, there has been substantial progress in understanding mechanisms which constrain the structure of the pedestal, and developing models to predict this structure. Here we employ the EPED1 pedestal model [2, 13, 14], which combines two physics constraints (calculated from sets of model equilibria that include a self-consistent treatment of the bootstrap current) to predict the pedestal height and width.

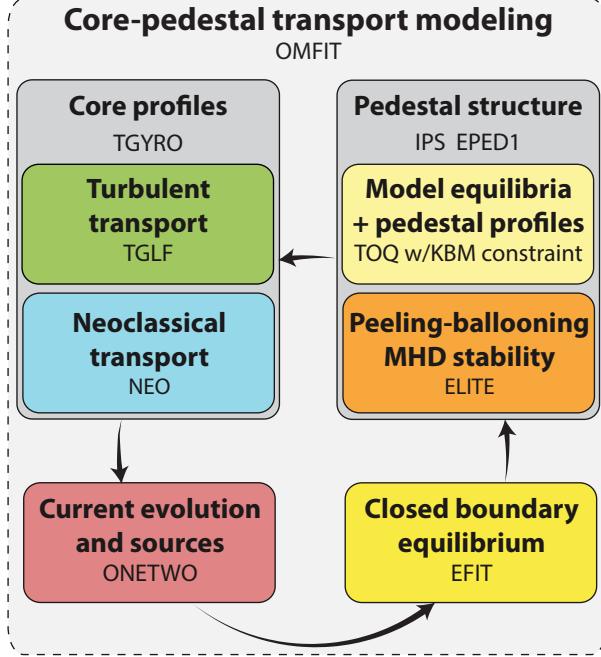


FIG. 3. Illustration of OMFIT integrated modeling workflow with dynamic pedestal. Coupled modules include core transport (turbulence and neoclassical), current and source evolution, magnetostatic equilibrium (EFIT), and EPED1 pedestal stability. The iterative workflow starts with the pedestal prediction from the EPED1 model. The profiles associated with the maximum pedestal height are used as an input to the ONETWO or TRANSP transport codes to calculate the sources of particles and heat in the plasma. The TGYRO code is then used to efficiently find the profiles for which the neoclassical fluxes from the NEO model and turbulent fluxes predicted by the TGLF model match those sources. The updated value of the global plasma pressure is used for a new run of the EPED1 model, and the cycle is repeated until the solution converges.

One important physics constraint is provided by peeling-ballooning (PB) modes, which are driven by the combination of strong pressure gradients, and resulting bootstrap current gradients, in the edge barrier region. PB modes are highly non-local, with a structure typically extending across the edge barrier into the outer core, and generally have toroidal mode numbers ranging from $\sim 4 - 30$ [33]. Special-purpose, efficient MHD codes, such as ELITE [33, 34] have been developed to calculate PB thresholds, and, employing diamagnetic corrections based on extended MHD, it has been possible to extensively compare the calculated PB threshold to observed constraints on the pedestal structure (and conditions for the onset of Type I edge localized modes (ELMs) or edge harmonic oscillations (EHO)), finding gener-

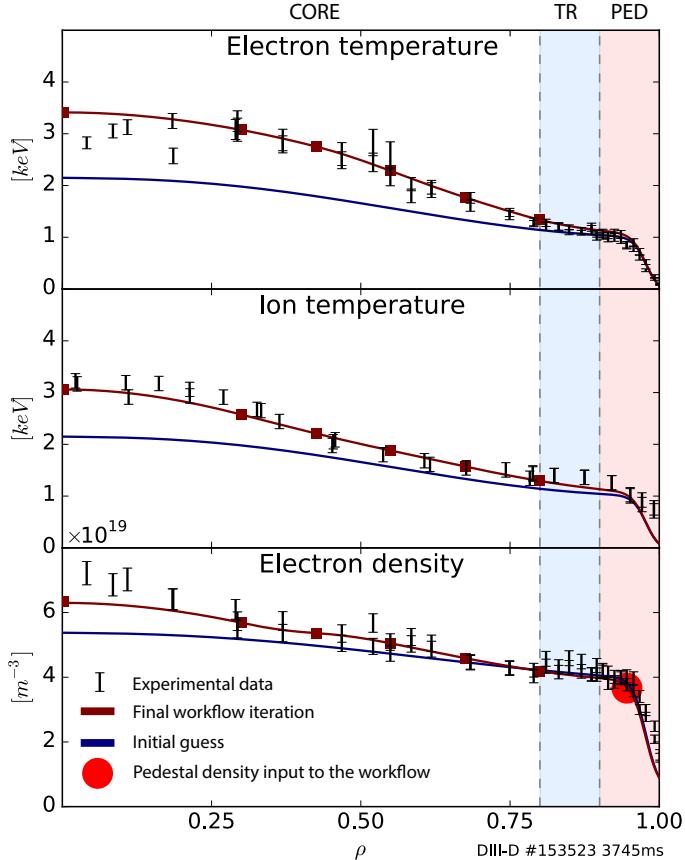


FIG. 4. Evolution of the simulated kinetic profiles for a DIII-D discharge and comparison with the experimental measurements. The profiles from the initial run of EPED are shown by the dark blue curves. The squares in these curves indicate the radii at which TGYRO performed the flux-matching calculations. The final step of the simulated profiles (dark red) are compared with experimental measurements, showing agreement across the entire plasma radius.

ally good agreement across large numbers of cases on several tokamaks [2, 35, 36]. However, the PB constraint by itself is not sufficient to self-consistently determine both the pedestal height and width, but rather gives a constraint on the height as a function of the width, which scales roughly with the $3/4$ power of the width [2].

While, strictly speaking, the scale separation between transport and macrostability breaks down in the pedestal, the EPED model proposes instead a separation between *nonlocal* modes (such as PB) and *nearly-local* modes which act primarily to constrain gradients within the electron transport barrier (ETB). We remark that the nonlocal modes are driven by free energy across the entire pedestal region, and can act to halt the penetration of

the edge transport barrier into the core – a process understood to be driven by strong diamagnetic $\mathbf{E} \times \mathbf{B}$ shear progressively shearing apart turbulent eddies near the pedestal top. EPED hypothesizes that, for high-performance (Type I ELM and Quiescent H-Mode) discharges, PB modes provide the global constraint. In general, there are many possible nearly-local modes, as well as significant levels of ion neoclassical energy transport, in the pedestal. EPED1 further conjectures that, because of the very strong $\mathbf{E} \times \mathbf{B}$ shear and generally strong particle sources in the ETB, the pressure gradient generally rises until it is finally constrained by the onset of the kinetic ballooning mode (KBM). This is proposed not because other modes do not exist or do not drive transport, but rather because they are not expected to provide a hard stability limit. More precisely, ion-scale modes, such as the ITG, are expected to be suppressed by the combination of $\mathbf{E} \times \mathbf{B}$ shear and β' stabilization. And typical electron-scale modes, such as the ETG, are driven by ratios of temperature to density gradients and are not generally expected to limit the pressure gradient by themselves in the presence of strong plasma sources.

The KBM threshold in the EPED model is calculated using the ballooning critical pedestal technique [13]. The KBM can be thought of as providing a constraint on the pressure gradient, or more quantitatively, the derivative of the local β_p (plasma pressure normalized to poloidal magnetic field pressure) with respect to the normalized poloidal flux, $d\beta_p/d\psi_N$. Because of the magnetic shear dependence of the KBM, together with the bootstrap current dependence on collisionality (which decreases strongly moving in from the separatrix), the KBM critical gradient generally increases moving radially inward. This implies that the radial average of the KBM critical gradient across the pedestal increases with the width of the pedestal (Δ), that is $\langle d\beta_p/d\psi_N \rangle \sim \Delta$, or, integrated across the pedestal, $\Delta = G\sqrt{\beta_{p,\text{ped}}}$, where $\beta_{p,\text{ped}}$ is the poloidal beta at the top of the pedestal (and it is assumed the poloidal beta at the separatrix is negligibly small) and G is a weakly varying function of collisionality and geometry such that $G \simeq 0.07 - 0.09$ (note that it is important to account for finite- n effects at low shear, which close off access to the *second stability* region that exists in high- n calculations). In the simplified EPED1 version of the EPED model employed here, G is fixed to a typical constant value, $G = 0.076$ [2, 13]. EPED1 then combines this condition for KBM criticality with a self-consistent calculation of the PB constraint with the ELITE code [1, 33, 34], using a series of roughly 100 model equilibria constructed using the TOQ [37] equilibrium solver, to determine the predicted pedestal height and width as the

intersection of the conditions for KBM and PB criticality (essentially solving two numerical equations for the two unknowns, pedestal height and width). The profiles in the *pedestal* region (highlighted red in Fig. 2) are given by EPED1, as a final result of roughly 700 PB calculations on 100 KBM-critical model equilibria.

The EPED1 model provides a prediction of the total pedestal pressure for a given electron density. In the pedestal and transition regions the electron and ions temperatures are assumed to be the same, and the ion density is calculated by enforcing quasi-neutrality and keeping their relative concentrations fixed. We note that there remain significant assumptions made in the coupled core-pedestal modeling done here. These include the supposition of sufficient power for operation in high performance H-mode. That is, we do not attempt to predict the conditions for the L-H transition or the transition into a high performance (Type I ELM or QH mode) regime. Furthermore, we are predicting approximate time-average conditions for the profiles, and are not explicitly accounting for transients such as ELMs.

B. Core-pedestal transition region: TR

The integrated modeling approach described here is new and differs significantly from the traditional methodology for predictive core transport simulations [3, 5–10, 12]. In traditional predictive modeling, a so-called *pivot radius* ρ_{pivot} is specified, at which the temperature and density are held fixed, and inside which the core transport solver evolves profiles. The pivot radius is nominally considered to represent the top of the pedestal and a typical value is $\rho_{\text{pivot}} \simeq 0.9$. In this approach, the pedestal is little more than a boundary condition for core modeling. More importantly, in the absence of a model for the pedestal, the choice of fixed density and temperature at the pivot radius is effectively arbitrary. Note that, in this approach, if ρ_{pivot} indeed signifies the top of the pedestal, then a transition region of zero radial extent between core and pedestal is implied. At the pedestal boundary however, the value of the inverse scale-length that satisfies the flux-matching constraint in TGYRO is unlikely to be consistent with the value that satisfies the PB and KBM constraints. When the disaccord between the two models is significant this manifests itself as an abrupt change in slope at the interface between the core and pedestal region (see for example the change in slope at the core-pedestal interface of the profiles in Fig. 15 of Ref. [9]; Figs. 16 and 17 of

Ref.[10]; Fig. 3, 5 and 6 of Ref. [12]). This is an important issue since, due to the stiffness of the core profiles, small changes in the gradients can lead to significant differences in the overall fusion performance. Recall that a transport model is said to be stiff if the transport fluxes increase rapidly with normalized temperature gradient, once the normalized profile gradient exceeds a threshold value.

If we let ρ_{core} denote the maximum radius for which TGYRO flux-matching procedure is valid, then what remains is a *transition region*

$$\rho_{\text{core}} < \rho < \rho_{\text{ped}} , \quad (19)$$

where ρ_{ped} is the minimum radius for which the EPED1 model can describe the profile (i.e. pedestal) structure. In the transition region, the H-mode kinetic profiles tend to form what could be described as a plateau. The physics processes that govern the profile evolution in the transition region changes from neoclassical and turbulent radial transport to PB and KBM stability constraints, including the effects of intermittent ELMs driven by moderate to high- n MHD instabilities. The transition region allows the gradient scale lengths that are consistent with the transport calculation in the core to smoothly transition to the values that are consistent with the PB and KBM dynamics of the pedestal.

IV. PROFILE PREDICTION WITH DYNAMIC EQUILIBRIUM AND PEDESTAL

The steady-state solution to the integrated simulation of pedestal structure, core transport, and magnetic equilibrium is found by means of the iterative workflow show in Fig. 3. This flow chart shows in more detail the individual components that are used for our particular implementation of the iterative workflow that was introduced earlier in Fig. 1B. In addition to the EPED1 model and the TGYRO transport solver used for the pedestal and core transport problem, we used the ONETWO transport code [38] for the calculation of the heat and particle sources, as well as for the evolution of the plasma current profile. Finally, the magnetic equilibrium calculation is performed with the EFIT equilibrium code [39] with the constrains to reproduce the pressure and current profiles that are output from TGYRO and ONETWO, as well as a user-prescribed plasma boundary. The separation of these processes into discrete modules can be justified based on timescale separation of the associated physics, and indeed these processes (for example, equilibrium and transport) have

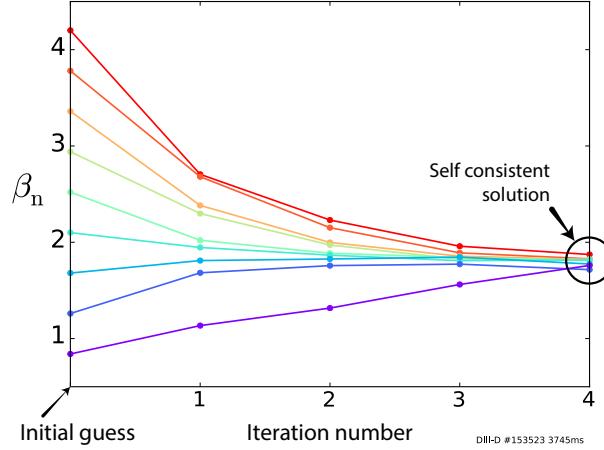


FIG. 5. Multiple integrated simulations of the same discharge were carried out for a broad range of initial guesses, β_n^{guess} , for the plasma pressure. The final pressure, β_n , converges to a result independent of the initial guess, thereby demonstrating robustness of the iterative method. Cases where the value of the initial guess is similar to the final solution evidently result in faster convergence, but this is not a requirement for the iteration scheme.

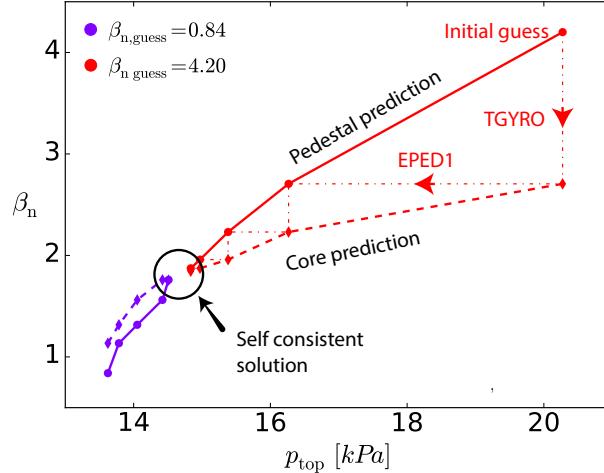


FIG. 6. Evolution of the coupled core-pedestal solution in the $\beta_n - p_{\text{top}}$ plane, for different initial guesses β_n^{guess} . Each update in the β_n solution (vertical shift) corresponds to a run of the core transport solver TGYRO, while each update in the pedestal height p_{top} (horizontal shift) corresponds to a run of the pedestal workflow EPED1. The graphical solution to the nonlinear coupled problem is constructed by connecting the points of the core and pedestal predictions. This figure shows that the solution is unique and independent of β_n^{guess} . This result demonstrates the robustness of the iterative scheme.

been treated separately, with success, for decades. Although this rules out the description of transient phenomena, in what follows we show that very close agreement with experimental results can be obtained with the present component coupling strategy.

The present workflow was implemented, and is available [40] via a coupling of the OMFIT [41] and Integrated Plasma Simulator (IPS) [42, 43] frameworks. Here the user controls the integrated simulation with OMFIT, which orchestrates the execution of the TGYRO, ONETWO, and EFIT components, while relying on the IPS to provide the High Performance Computing (HPC) enabled IPS-EPED1 workflow. OMFIT further handles data exchange across distributed workstations and clusters, as well as selected post-processing tasks. For the simulations presented herein, we required that the original EPED1 workflow [2] (yet another nested workflow) executes sufficiently rapidly so as to enable parameter sweeping of the entire coupled model in a few minutes. This was not practical with the original IDL based EPED1 workflow running on a local cluster. To achieve this goal we translated EPED1 into the IPS-EPED1 workflow, which can now take full advantage of HPC resources by allowing the IPS to schedule and launch the individual execution of the components internal to EPED1. This IPS-EPED1 workflow replicates the original EPED1 capability, but at a fraction of the wall clock time. As indicated in Fig. 3, the EPED1 model [2] maps the equilibria-stability parameter space using the TOQ and ELITE component codes. The IPS-EPED1 nested workflow utilizes the capabilities of the IPS framework to launch these tasks in parallel. A typical IPS-EPED1 iteration uses up to 700 cores on the Edison Cray XC30 platform at NERSC [44], and simulation results are obtained in less than 2 minutes (dominated by the time for one TOQ plus one ELITE run for the highest mode number). The IPS-EPED1 workflow has been verified against the original EPED1 for the database of 362 runs presented in [2]. The average difference for the calculated pedestal pressure between original EPED1 and IPS-EPED1 is 0.147 kP, which is within the algorithmic uncertainty of EPED. This verification scan was conducted using the integrated IPS-DAKOTA optimization and parameter sweep environment [45]. Using this environment, multiple IPS-EPED1 executions are dispatched concurrently within a single instance of the IPS, with DAKOTA coordinating the parameterization of each run and aggregating the final results.

In this section we apply the iterative workflow to simulation of a DIII-D discharge. The discharge – chosen because of its relevance to the ITER baseline scenario – is characterized by low torque and dominant electron heating. Simulation results are summarized in Fig. 4.

The inputs to the simulation are the plasma shape, B_t , I_p , the configuration of the heat, particle and current sources (neutral beam and electron cyclotron radio-frequency heating), the pedestal electron density $n_{e,\text{ped}}$ (indicated by a red circle in Fig. 4) and the ion effective charge Z_{eff} . For the core region, we set the innermost TGYRO flux-matching radius at $\rho_1 = 0.3$ (i.e., the sawtooth inversion radius) and the last at $\rho_{\text{core}} = 0.8$. The transition region, next, begins at $\rho_{\text{core}} = 0.8$ and ends at $\rho_{\text{ped}} = 0.9$. The pedestal region, finally, covers the region $\rho_{\text{ped}} < \rho < 1$. To illustrate the predictive capabilities of this workflow, the initial value of the normalized plasma beta was chosen to be $\beta_n^{\text{guess}} = 1.26$, a value that is intentionally far from the experimental value of $\beta_{n,\text{exp}} = 1.69$.

The iterative workflow starts with the pedestal prediction from the EPED1 model. Internally the EPED1 model generates a series of KBM-critical equilibria and profiles with increasing pedestal pressure to find the maximum pedestal height p_{top} and width w_{top} consistent with PB stability. When PB criticality is reached it is expected that an ELM or edge harmonic oscillation will be triggered, preventing further increase of the height and width) [2, 13]. The profiles associated with the maximum pedestal height are used as an input to the ONETWO transport code to calculate the sources of particles and heat in the plasma. The TGYRO code the computes the inverse scale-lengths for which the neoclassical and turbulent fluxes (as predicted by the NEO and TGLF models) match those sources. Smooth temperature and density profiles (at any desired number of flux-matching radii) for both electron and ion are then evaluated using the integrals defined by Eq. 16. The value of β_n from the updated equilibrium calculation is used for a new run of the EPED1 model. This cycle is repeated until the solution converges. The final iteration of the simulated profiles (dark red) is in good agreement with experimental measurements, and the final value of the normalized plasma beta $\beta_n^{\text{final}} = 1.68$ closely matches the experimental value of 1.69. Note that the final converged solution is independent of the initial guess β_n^{guess} , as illustrated in Fig. 5. In this plot, multiple self-consistent simulations of the same discharge were carried out for a broad range of β_n^{guess} , yielding the important result that all converge to the same steady-state solution. Cases where the value of the initial guess is similar to the final solution will of course result in faster convergence, but a very accurate initial guess is evidently not a requirement for convergence of the iteration scheme.

A deeper understanding of the robustness of the core-pedestal coupling approach can be obtained by visualizing the evolution of the solution in the $\beta_n - p_{\text{top}}$ plane, as shown in Fig. 6.

We define p_{top} as the value of the pressure profile evaluated at 1.5 times the pedestal width in from the separatrix. This plane is a natural choice for studying the coupled core-pedestal problem since β_n is modified by TGYRO for fixed p_{top} , whereas EPED1 modifies p_{top} at fixed β_n . The circles and diamonds in Fig. 6 represent the value of the global pressure at the top of the pedestal after the execution of the pedestal and core solvers, respectively. The dash-dotted segments mark the progression of the solution through the execution of the two solvers at successive iterations for $\beta_n^{\text{guess}} = 4.20$. The figure also illustrates convergence starting at a much lower value of $\beta_n^{\text{guess}} = 0.84$. The graphical solution to the nonlinear coupled problem is thus constructed by connecting the points of the core and pedestal predictions. Values of β_n above the core prediction curve are precluded by the effect of stiff core transport, while values of p_{top} higher than the pedestal prediction curve are precluded by the onset of PB modes (given KBM-critical profiles). This analysis provides confidence that the iterative scheme is stable and that a unique solution can be found for the set of parameters that are input to the workflow.

V. OPTIMIZATION OF ITER BASELINE FUSION PERFORMANCE

A self-consistent optimization of the fusion power of the ITER baseline scenario was carried out based on the iterative workflow described in the previous section. The input parameters to the workflow were taken from the ITER baseline case studied in Ref. [10] (see Fig. 6 of Ref. [10]). In order to use the same input parameters as in Ref. [10], we modified the iteration scheme of Fig. 3 to use the TRANSP [46] code for the evolution of the current profile, the evaluation of the heating sources, and calculation of the Grad-Shafranov magnetic equilibrium. For this study we varied the pedestal density $n_{e,\text{ped}}$ and effective ion charge $Z_{\text{eff,ped}}$, as input to the EPED1 model. A total of 16 simulations were performed, with four variations both in the pedestal density and $Z_{\text{eff,ped}}$. Figure 7 shows four simulations with $Z_{\text{eff,ped}} = 1.7$ (the lowest value), and Fig. 8 shows four simulations with $Z_{\text{eff,ped}} = 3.4$ (the highest value). In both cases we plot the predicted pressure, electron and ion temperatures, and electron densities. We emphasize that only Z_{eff} in the pedestal is varied, while the core value is kept constant among all cases (See Fig. 13 of Ref. [10]). Work to include self-consistent impurity sources and transport is underway as part of the AToM SciDAC project.

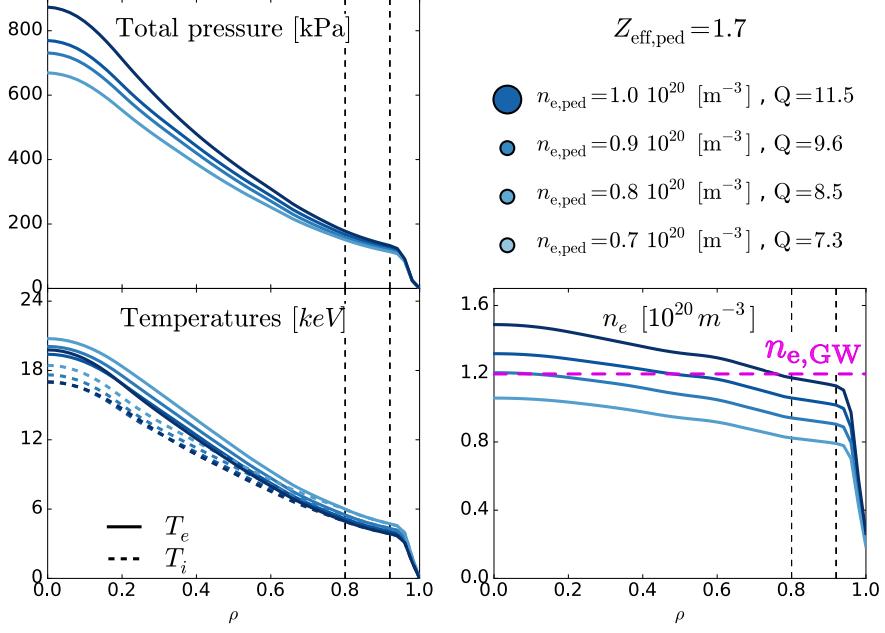


FIG. 7. Results of a self-consistent optimization study for the ITER baseline scenario for $Z_{\text{eff,ped}} = 1.7$ illustrating plasma performance for four values of the pedestal density, $n_{\text{e,ped}}$. Here fusion performance improves with increased pedestal density, which confirms naive expectations.

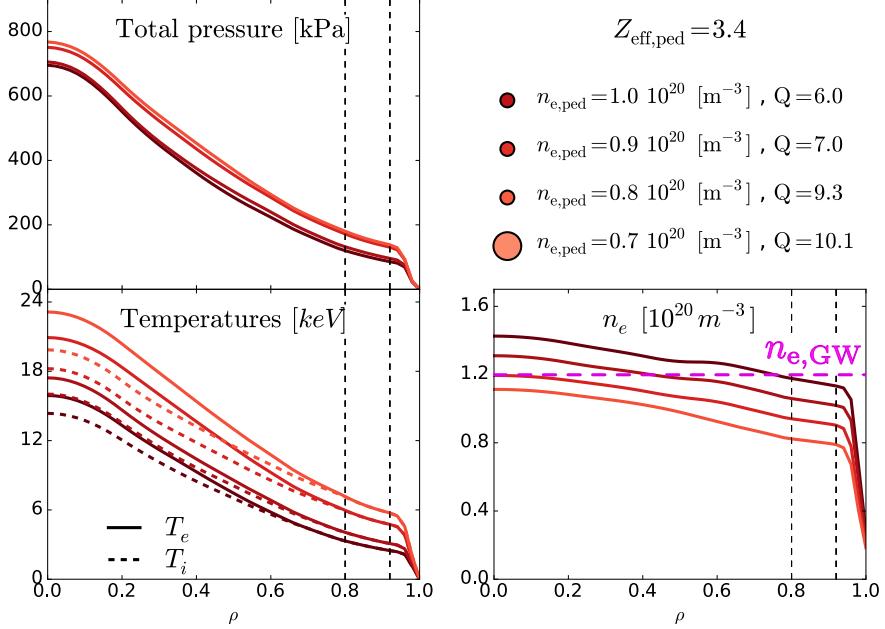


FIG. 8. Same as Fig. 7, except for $Z_{\text{eff,ped}} = 3.4$. However, in contrast to Fig. 7, fusion performance now **decreases** with increasing pedestal density, illustrating the critical importance of a self-consistent pedestal model for ITER modeling.

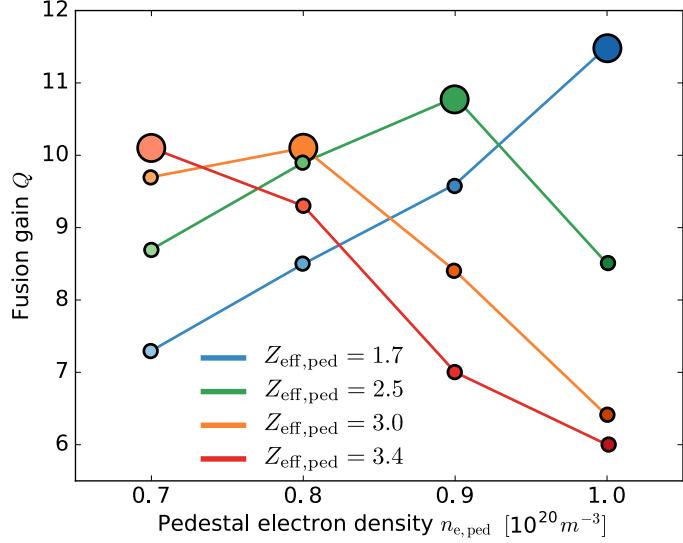


FIG. 9. Illustration of fusion gain Q as a function of n_e for different values of the pedestal Z_{eff} . The counter-intuitive trend in n_e for large Z_{eff} is apparent.

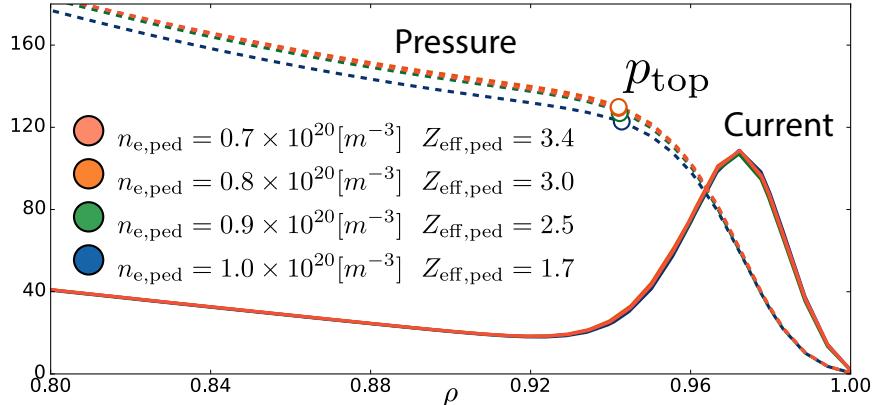


FIG. 10. Although the best performing cases of the ITER baseline optimization study occur for different values of $n_{e,\text{ped}}$ and $Z_{\text{eff},\text{ped}}$, they exhibit nearly identical pedestal pressure and current profile. As such, their edge MHD stability properties are also comparable.

For each of the panels, the best performing cases are highlighted with a larger marker symbol in the legend. As a consequence of the strong core transport stiffness, the core profiles are self-similar with respect to a rigid translation of the boundary condition at the interface between the transition region and the core region. In all cases the density exhibits a profile peaking of about $n_{e_0}/\langle n_e \rangle \sim 1.3$. For low values of the effective ion charge ($Z_{\text{eff},\text{ped}} = 1.7$) the fusion gain Q increases with the pedestal density, which is the naive expectation. However,

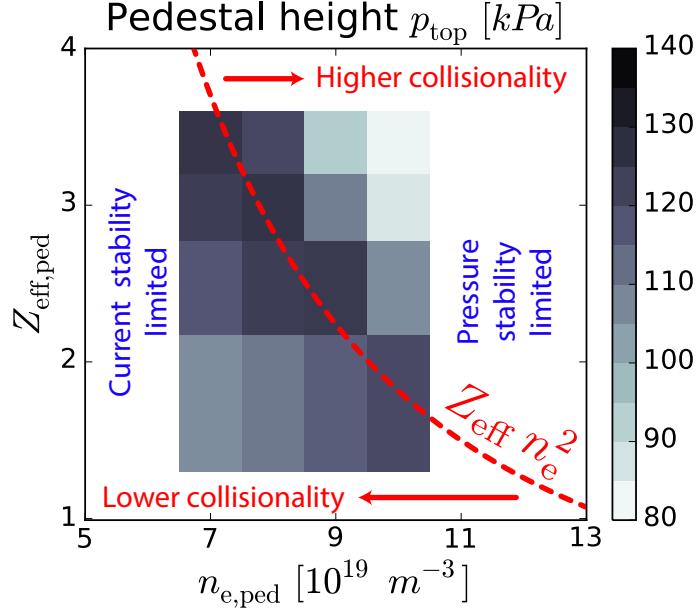


FIG. 11. The zeroth-order scaling of the bootstrap current with collisionality at fixed pressure is such that $j_{BS} \propto Z_{\text{eff}} n_e^2$. Such scaling is consistent with the observed similar current profiles for the best performing cases in Fig. 10.

in this case the highest fusion performance exceeds the empirical Greenwald density limit [47], which for the ITER baseline scenario (with plasma current of 15 MA and a minor radius of 2 m) is approximately equal to $n_{GW} \sim 1.2 \times 10^{20} \text{ m}^{-3}$. As $Z_{\text{eff,ped}}$ is increased from 1.7 to 3.4, the pedestal density at which the peak fusion power is reached occurs for lower values of $n_{e,\text{ped}}$. And, counterintuitively, for $Z_{\text{eff,ped}} = 3.4$, the fusion performance decreases with increasing pedestal density, as shown in Fig. 8. More precisely, for $Z_{\text{eff,ped}} = 1.7, 2.5, 3.0, 3.4$, the optimal pedestal density is $n_{e,\text{ped}}[10^{20}/\text{m}^3] = 1.0, 0.9, 0.8, 0.7$. This fusion performance trend is further illustrated in Fig. 9.

Although the best performing cases for the ITER baseline optimization study occur for different values of $n_{e,\text{ped}}$ and $Z_{\text{eff,ped}}$, they share similar pedestal pressure and current profiles (see Fig. 10). As such, their edge PB and KBM properties are also comparable. As a consequence of the core transport stiffness, the core pressure (and thus the fusion performance) closely correlates with the pedestal height. This trend is well-captured by the grayscale pedestal height map of Fig. 11. As expected by the zeroth-order scaling of the bootstrap current with collisionality at fixed pressure $j_{BS} \propto Z_{\text{eff}} n_e^2$, the highest values of the pedestal height are aligned with contours of $Z_{\text{eff}} n_e^2$. This explains why the best performing

cases shown in Fig. 10 have similar current profiles. Based on the collisional scaling of the bootstrap current, the loss of performance at lower and higher collisionality is attributed to the effect of primarily current and pressure-driven PB modes, respectively.

VI. CONCLUSIONS AND FUTURE PERSPECTIVES

In the work we have given an overview of a fundamentally new integrated modeling capability which includes, among other key physics modules, a *dynamic pedestal* component that can be used to ensure the consistency of core and pedestal evolution. More generally, we have described an approach to modeling that takes into account the strong interplay between (a) core transport, (b) pedestal structure, (c) current profile evolution and (d) magnetic equilibrium. This capability improves the accuracy and reliability of previous approaches to the prediction of fusion performance. In this iterative workflow, self-consistency is achieved by successive iterations of pedestal model, core model, sources and equilibrium calculation. We have shown that the coupling scheme can closely reproduce the experimental measurements *without* requiring a priori knowledge of the temperature and density profiles. Further, we have demonstrated that the self-consistent, iterative solution is independent of the initial guess that is used to initiate the pedestal calculation at the beginning of the iterative process.

Results of a self-consistent optimization of the ITER baseline scenario show that $n_{e,ped}$ as well as $Z_{eff,ped}$ are strong actuators for the fusion power, both of which will likely need to be actively controlled during ITER operations to optimize the fusion performance, and satisfy the requirements imposed by the density limit. The dependency with respect to $Z_{eff,ped}$ points to the need for coupling to a Scrape-Off-Layer (SOL) in the iterative workflow: an important future development which is underway as part of the AToM SciDAC project. The TGYRO transport code will also be updated to independently evolve the kinetic profile of the individual ion species.

The coupling workflow described in this document has been implemented within the OMFIT integrated modeling framework. As such, this tool is available to the broader fusion community, to further validate the individual physics components of the workflow as well as their couplings, and enable the self-consistent design of steady-state operational scenarios from first-principles calculation.

VII. ACKNOWLEDGMENTS

Work supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-SC0012656 (GA AToM SciDAC), DE-AC05-00OR22725 (ORNL AToM SciDAC), DE-SC0012633 (UCSD AToM SciDAC), DE-FG02-95ER54309 (GA theory), DE-FC02-06ER54873 (ESL), DE-FC02-04ER54698 (DIII-D). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

- [1] P. Snyder, K. Burrell, H. Wilson, M. Chu, M. Fenstermacher, A. Leonard, R. Moyer, T. Osborne, M. Umansky, W. West, *et al.*, Nuclear Fusion **47**, 961 (2007).
- [2] P. Snyder, R. Groebner, A. Leonard, T. Osborne, and H. Wilson, Physics of Plasmas **16**, 056118 (2009).
- [3] J. Kinsey, G. Bateman, T. Onjun, A. Kritz, A. Pankin, G. Staebler, and R. Waltz, Nucl. Fusion **43**, 1845 (2003).
- [4] J. Kinsey, G. Staebler, and R. Waltz, Phys. Plasmas **15**, 055908 (2008).
- [5] J. Candy, C. Holland, R. Waltz, M. Fahey, and E. Belli, Phys. Plasmas **16**, 060704 (2009).
- [6] J. Kinsey, G. Staebler, J. Candy, R. Waltz, and R. Budny, Nucl. Fusion **51**, 083001 (2011).
- [7] M. Murakami, J. Park, G. Giruzzi, J. Garcia, P. Bonoli, R. Budny, E. Doyle, A. Fukuyama, N. Hayashi, M. Honda, A. Hubbard, S. Ide, F. Imbeaux, E. Jaeger, T. Luce, Y.-S. Na, T. Oikawa, T. Osborne, V. Parail, A. Polevoi, R. Prater, A. Sips, J. Snipes, H. S. John, P. Snyder, I. Voitsekhovitch, and I. O. S. Group, Nuclear Fusion **51**, 103006 (2011).
- [8] G. Bateman, T. Onjun, and A. H. Kritz, Plasma Physics and Controlled Fusion **45**, 1939 (2003).
- [9] R. Budny, R. Andre, G. Bateman, F. Halpern, C. Kessel, A. Kritz, and D. McCune, Nucl. Fusion **48**, 075005 (2008).
- [10] R. Budny, Nuclear Fusion **49**, 085008 (2009).
- [11] M. Murakami, J. Park, D. Batchelor, S. Diem, W. Elwasif, A. Sontag, and D.-D. Team, in *APS Meeting Abstracts*, Vol. 1 (2013) p. 8021P.

- [12] F. Poli, C. Kessel, P. Bonoli, D. Batchelor, R. Harvey, and P. Snyder, Nuclear Fusion **54**, 073007 (2014).
- [13] P. Snyder, R. Groebner, J. Hughes, T. Osborne, M. Beurskens, A. Leonard, H. Wilson, and X. Xu, Nuclear Fusion **51**, 103016 (2011).
- [14] P. Snyder, T. Osborne, K. Burrell, R. Groebner, A. Leonard, R. Nazikian, D. Orlov, O. Schmitz, M. Wade, and H. Wilson, Phys. Plasmas **19**, 056115 (2012).
- [15] ITER team and others, Nucl. Fusion **39**, 2137 (1999).
- [16] AToM, *SciDAC project*, <http://scidac.github.io/atom/>.
- [17] G. Staebler, J. Kinsey, and R. Waltz, Phys. Plasmas **14**, 055909 (2007).
- [18] J. Candy and R. Waltz, J. Comput. Phys. **186**, 545 (2003).
- [19] E. Belli and J. Candy, Plasma Phys. Control. Fusion **50**, 095010 (2008).
- [20] E. Belli and J. Candy, Plasma Phys. Control. Fusion **51**, 075018 (2009).
- [21] E. Belli and J. Candy, Phys. Plasmas **17**, 112314 (2010).
- [22] T. Rhodes, C. Holland, S. Smith, A. White, K. Burrell, J. Candy, J. DeBoo, E. Doyle, J. Hillesheim, J. Kinsey, G. McKee, D. Mikkelsen, W. Peebles, C. Petty, R. Prater, S. Parker, Y. Chen, L. Schmitz, G. Staebler, R. Waltz, G. Wang, Z. Yan, and L. Zeng, Nuclear Fusion **51**, 063022 (2011).
- [23] C. Holland, C. Petty, L. Schmitz, K. Burrell, G. McKee, T. Rhodes, and J. Candy, Nuclear Fusion **52**, 114007 (2012).
- [24] J. Citrin, F. Jenko, P. Mantica, D. Told, C. Bourdelle, R. Dumont, J. Garcia, J. Haverkort, G. Hogewij, T. Johnson, M. Pueschel, and J.-E. contributors, Nuclear Fusion **54**, 023008 (2014).
- [25] A. B. Navarro, T. Happel, T. Grler, F. Jenko, J. Abiteboul, A. Bustos, H. Doerk, D. Told, and A. U. Team, Physics of Plasmas **22**, 042513 (2015), <http://dx.doi.org/10.1063/1.4919022>.
- [26] A. E. White, N. T. Howard, A. J. Creely, M. A. Chilenski, M. Greenwald, A. E. Hubbard, J. W. Hughes, E. Marmar, J. E. Rice, J. M. Sierchio, C. Sung, J. R. Walk, D. G. Whyte, D. R. Mikkelsen, E. M. Edlund, C. Kung, C. Holland, J. Candy, C. C. Petty, M. L. Reinke, and C. Theiler, Physics of Plasmas **22**, 056109 (2015), <http://dx.doi.org/10.1063/1.4921150>.
- [27] W. Houlberg, K. Shaing, S. Hirshman, and M. Zarnstorff, Phys. Plasmas **4**, 3230 (1997).
- [28] E. Belli and J. Candy, Plasma Phys. Control. Fusion **54**, 015015 (2012).
- [29] G. Bateman, A. Kritz, J. Kinsey, A. Redd, and J. Weiland, Phys. Plasmas **5**, 1793 (1998).

[30] R. Waltz, G. Staebler, W. Dorland, G. Hammett, M. Kotschenreuther, and J. Konings, *Phys. Plasmas* **4**, 2482 (1997).

[31] G. Staebler, J. Kinsey, and R. Waltz, *Phys. Plasmas* **12**, 102508 (2005).

[32] J. Lee, M. Barnes, F. Parra, E. Belli, and J. Candy, *Plasma Phys. Control. Fusion* **57**, 125006 (2015).

[33] P. Snyder, H. Wilson, J. Ferron, L. Lao, A. Leonard, T. Osborne, and A. Turnbull, *Physics of Plasmas* **9**, 2037 (2002).

[34] H. Wilson, P. Snyder, and G. Huysmans, *Phys. Plasmas* **9**, 1277 (2002).

[35] T. H. Osborne, P. B. Snyder, K. H. Burrell, T. E. Evans, M. E. Fenstermacher, A. W. Leonard, R. A. Moyer, M. J. Schaffer, and W. P. West, *Journal of Physics: Conference Series* **123**, 012014 (2008).

[36] S. Saarelma and A. Alfier and M. N. A. Beurskens and R. Coelho and H. R. Koslowski and Y. Liang and I Nunes and JET EFDA Contributors, *Plasma Phys. Control. Fusion* **51**, 035001 (2009).

[37] L. L. Lao, S. P. Hirshman, and R. M. Wieland, *Physics of Fluids* **24**, 1431 (1981).

[38] H. St John, T. Taylor, Y. Lin-Liu, and A. Turnbull, in *Presented at the 15th International Conference on Plasma Physics and Controlled Nuclear Fusion Research, Seville, Spain, 26 Sep.-1 Oct. 1994*, Vol. 1 (1994).

[39] L. Lao, J. Ferron, R. Groebner, W. Howl, H. S. John, E. Strait, and T. Taylor, *Nuclear Fusion* **30**, 1035 (1990).

[40] OMFIT, *One Modeling Framework for integrated Tasks*, <http://gafusion.github.io/OMFIT-source/>.

[41] O. Meneghini, S. Smith, L. Lao, O. Izacard, Q. Ren, J. Park, J. Candy, Z. Wang, C. Luna, V. Izzo, B. Grierson, P. Snyder, C. Holland, J. Penna, G. Lu, P. Raum, A. McCubbin, D. Orlov, E. Belli, N. Ferraro, R. Prater, T. Osborne, A. Turnbull, and G. Staebler, *Nuclear Fusion* **55**, 083008 (2015).

[42] W. Elwasif, D. Bernholdt, A. Shet, S. Foley, R. Bramley, D. Batchelor, and L. Berry, in *2010 18th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)* (IEEE, 2010) pp. 419–427.

[43] W. Elwasif, D. Bernholdt, S. Foley, A. Shet, and R. Bramley, in *Computer Systems and Applications (AICCSA), 2011 9th IEEE/ACS International Conference on* (2011) pp. 188–

195.

- [44] NERSC, *National Energy Research Scientific Computing Center* , <https://www.nersc.gov>.
- [45] W. Elwasif, D. Bernholdt, S. Pannala, S. Allu, and S. Foley, in *Computational Science and Engineering (CSE), 2012 IEEE 15th International Conference on* (2012) pp. 102–110.
- [46] TRANSP, *integrated modeling code*, <http://w3.ipppl.gov/transp/>.
- [47] M. Greenwald, J. Terry, S. Wolfe, S. Ejima, M. Bell, S. Kaye, and G. Neilson, Nuclear Fusion **28**, 2199 (1988).