J. Plasma Phys. (2026), *vol.* 92, E67 © The Author(s), 2026. Published by Cambridge University Press.
This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (https://creativecommons.org/licenses/by-nc-nd/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use and/or adaptation of the article.
doi:10.1017/S0022377826101573

Performance and transport in the ARC tokamak

N.T. Howard[1], P. Rodriguez-Fernandez[1], J. Hall[1], M. Muraca[1], A. Saltzman[1], A. Ho[1], J.C. Hillesheim[2], A.J. Creely[2], T.H. Eich[2], T. Body[2], P.B. Snyder[2] and C. Holland[3]

[1]Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[2]Commonwealth Fusion Systems, Devens, MA 01434, USA
[3]University of California-San Diego, La Jolla, CA 92093, USA

Corresponding author: N.T. Howard, nthoward@psfc.mit.edu

(Received 17 October 2025; revision received 16 January 2026; accepted 20 February 2026)

The ARC^{TM} tokamak, a high-field ($B_T = 11.4$ T) fusion power plant, under development by Commonwealth Fusion Systems, is studied using a suite of integrated modelling tools to predict its fusion power generation (P_{fus}), transport and confinement properties. Analysis is based off an ARC operational point scoped first with zero-dimensional (0-D) plasma operational contour (POPCON) modelling to produce 1.13 GW of fusion power. A suite of integrated modelling tools (TRANSP, ASTRA and TORAX) were applied to predict the performance and kinetic profiles of the ARC design point, yielding a range of predicted performance spanning from ∼900 to 1300 MW in rough quantitative agreement with POPCON predictions. The sensitivity of these results to uncertain modelling inputs was probed using scans of pedestal boundary conditions around EPED-predicted values (total pressure and temperature ratios), tungsten concentration and seperatrix density around their nominal assumptions. Pedestal pressure and pedestal top (T_i/T_e) play a large role in 1.5-dimensional performance predictions, able to modify the predicted P_{fus} by a factor of 2 within reasonable assumptions. High-fidelity core nonlinear gyrokinetic profile predictions, performed using CGYRO (Candy *et al.* 2016 *J. Comput. Phys.*, vol. 324, pp. 73–93) coupled with the PORTALS (Rodriguez-Fernandez *et al.* 2024 *Nucl. Fusion*, vol. 64, 076034; *Phys. Plasmas*, vol. 31, 2024, 062501) framework, yield substantially lower performance ($P_{fus} = 677$ MW) compared with 0-D and medium-fidelity modelling for nominal assumptions, showing that there is non-negligible uncertainty between models and that future work on SPARC may help resolve discrepancies. Lower overall performance results from significantly reduced volume-averaged densities and temperatures, along with reduced levels of density and temperature peaking. Turbulence and transport are largely dominated by ion temperature gradient across the profile, confirmed by both linear stability and the response of the nonlinear fluxes to changes in gradients, with some impact of kinetic ballooning modes in the deep core. This work represents one of the most complete scoping of potential fusion power plant conditions performed to date. The extensive integrated modelling provides confidence in ARC performance approaching 1 GW, while nonlinear gyrokinetic modelling results in open questions

SPARC® and ARC^{TM} are trademarks of Commonwealth Fusion Systems®.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

into the physics of density and temperature peaking in fusion-power-plant-relevant operational space. A discussion of results and the role that the SPARC tokamak (Creely *et al.* 2020 *J. Plasma Phys.*, vol. 86, 865860502) will play in informing ARC design, performance and operation is presented.

Key words: plasma confinement, fusion plasma

1. Introduction

Recent worldwide emphasis on achieving the goal of low-cost, carbon-free electricity generation has provided renewed momentum for the development of energy production via fusion. In addition to investments from public entities worldwide, a vibrant fusion industry ecosystem has evolved to play a crucial role in pushing beyond fusion from science and technology research to the development of commercial fusion power plants (FPPs). Commonwealth Fusion Systems (CFS) is a private fusion company headquartered in Devens, MA which is pursuing a high-field path to fusion energy based on the tokamak concept. In 2021 CFS completed the construction and testing of the world's first high-field tokamak superconducting magnet, the Toroidal Field Model Coil, to demonstrate that high-temperature superconductors offered a viable path to creating fusion devices at high magnetic field (>12 T). CFS's proposed development of tokamak-based FPPs leverages years of existing physics-based work while reducing the overall capital cost for fusion devices by utilising more compact designs enabled by high-temperature superconductor magnets. Currently, the SPARC tokamak (Creely *et al.* 2020) is under construction by CFS to achieve $Q > 1$ in a compact high-field device. However, SPARC is only an intermediate step between current devices and a tokamak capable of generating electricity from fusion. The next device along the fusion roadmap is the development of the ARC tokamak, a high-field FPP targeting the production of gigawatt levels of fusion power.

As we transition from current fusion devices, focused on scientific discovery, to those capable of creating electricity from fusion, the need for accurate scoping of the performance in these devices has also increased – to aid in design, optimisation and eventual safe operation. Historically, tokamak designs such as ITER were based purely on empirical scaling laws derived from multi-machine databases of observational data, with more rigorous modelling only occurring after the initial machine design. However, the advent of more reliable models for plasma transport, heating and equilibrium has made physics-based projections of tokamak performance possible, and routine enough to be used for machine design. Most such studies have relied on reduced, physics-based transport models, which are fast and relatively accurate, but are generally 'tuned' to reproduce the results of more first-principles-based models such as nonlinear gyrokinetics. Nearly 30 years of validation efforts on tokamaks worldwide (Ross & Dorland 2002; Candy & Waltz 2003; White *et al.* 2008; Casati *et al.* 2009; Holland *et al.* 2009; Angioni *et al.* 2011; Howard *et al.* 2012; Casson *et al.* 2013; Guttenfelder *et al.* 2013; White *et al.* 2013; Citrin *et al.* 2014 Field *et al.* 2014; Gorler *et al.* 2014; Navarro *et al.* 2015; Bonanomi *et al.* 2018) have provided confidence that nonlinear gyrokinetics is in fact an accurate model of turbulence and transport in the tokamak core, but computational cost has made it generally prohibitively expensive for profile predictions. However, utilising high-performance computing, machine learning techniques and GPU acceleration, it is now possible to perform predictions of fusion devices at high fidelity, enabling more accurate

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

modelling of fusion plasmas. Throughout this paper we refer to both 'medium'- and 'high'-fidelity integrated modelling. For the purposes of this work medium fidelity refers to predictions based on physics-based models which include quasi-linear transport models such as TGLF (Staebler, Kinsey & Waltz 2007) and Qualikiz (Citrin *et al.* 2017) which capture crucial physics relevant for core plasma turbulence but rely on saturation rules determined from nonlinear gyrokinetics. In contrast, high fidelity refers to profile predictions that utilise nonlinear gyrokinetics to perform profile predictions.

In this paper we (i) examine predictions of fusion power and gain (Q) of the ARC V3A design point, (ii) determine how uncertain input parameters affect the overall accuracy of our projections and (iii) obtain insight into the turbulence and transport properties of ARC. This analysis focused on the most current ARC design (V3A), but it is anticipated that the design will continue to evolve and that final ARC designs will leverage new insights gained through operation of SPARC. The remainder of this paper is organised as follows. Section 2 provides an overview of the ARC tokamak for the reader, a description of the performance goals and discussion of the operational point examined in this paper. Section 3 focuses on physics-based, medium-fidelity modelling. We give a brief overview of the analysis codes that were utilised to produce integrated modelling predictions of the ARC operational point and compare the predictions of different models with each other and zero-dimensional (0-D) approaches. The sensitivity of results to poorly constrained modelling inputs is examined and the implications discussed. Section 4 presents results focused on high-fidelity, nonlinear gyrokinetic predictions of ARC performance and transport. Profile predictions based on neoclassical and gyrokinetic turbulence modelling are presented and contrasted with lower-fidelity modelling. We discuss the nature of the turbulence in ARC and the implications of this turbulence for accurate modelling. Finally, § 5 provides a brief overview of this paper's findings, discusses the sensitivity of modelling results and discusses important avenues for research on existing devices that can close open questions for FPPs such as ARC.

2. ARC performance goals, operational point and POPCON modelling

The ARC tokamak is currently in development by CFS. This device is slated for construction and operations in the 2030 s. ARC is part of the high-field tokamak path to fusion energy and the natural progression of the SPARC tokamak (Creely *et al.* 2020; Rodriguez-Fernandez, Howard & Candy 2022) currently under construction in Devens, MA. The origin of the ARC tokamak concept arises from a paper published by Sorbom *et al.* (2015). However, significant evolution in this concept design point has occurred since the original paper. In fact, the 2015 paper featured only 0-D evaluation of the core performance and a comparatively low-fidelity investigation of the magnets and structural design. The version of ARC presented here results from a more rigorous investigation of all aspects of the power plant design point. ARC is a high-field tokamak that utilises high-temperature superconducting magnets, anticipated to leverage REBCO technology (Seeber 2023) currently being used in the construction of the SPARC tokamak. ARC V3A is a large tokamak with major and minor radii of $R = 4.62$ m, $a = 1.18$ m, yielding a plasma volume of 189 m^3, making this device larger than the JET tokamak ($R = 2.96$ m) but smaller than ITER ($R = 6.2$ m). In contrast to the original paper, ARC is expected to use pulsed operation with estimated flattop times of 900 s and will operate with an on-axis toroidal field $B_T = 11.4$ T and a plasma current of $I_p = 12.0$ MA, resulting in

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

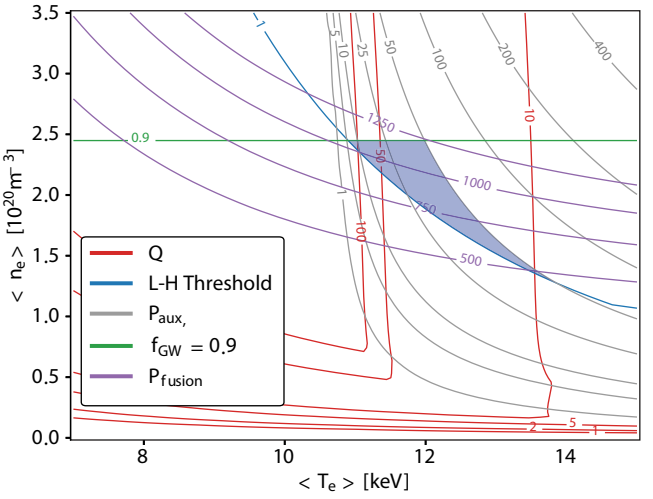


FIGURE 1. The POPCON for the ARC design point is plotted. The shaded blue region indicates the currently projected operational window which is subject to future design updates.

$q_{95} = 3.7$ and $q_* = 3.0$. A flexible coil set is expected to enable significant shaping, providing access to higher-performance operational regimes with seperatrix elongations and triangularity of $\kappa_{sep} = 1.8$ and $\delta_{sep} = 0.65$. Up to 50 MW of auxiliary heating will be available, utilising H-minority heating via ion cyclotron resonance heating (ICRH) at ~180 MHz to provide on-axis heating. The results covered in this paper are indicative of conditions that should be obtained shortly after the achievement of the current flattop where current profile is expected to be mostly relaxed, but small changes in the evolution of the profile may continue to occur within the 900 s flattop. Changes to the current profile (and safety factor profile) may have non-negligible impact on predicted performance. However, such investigations are out of the scope of this paper and are left for future work. In this paper we therefore assume an approximate steady-state solution of plasma current density profile. It is also worth noting that the ARC design assumes the use of demountable coils. This provides a unique flexibility to modify and optimise the vacuum vessel to provide access to slightly different plasma shapes and volumes – with different vessels representing a trade-off between performance and component lifetimes due to necessary adjustment to shielding. For a more complete description of the ARC tokamak, including descriptions of coils and vessel components, we refer the reader to Hillesheim *et al*. (2026). Two such configurations were investigated as part of this work and are described briefly in § 3.

The ARC point was chosen to provide approximately 1.13 GW of total fusion power generation with an anticipated ~415 MW of net electricity, after efficiencies and subtracting the power to run the plant. The operational point was designed using the well-known plasma operational contour (POPCON) modelling technique (Houlberg, Attenberger & Hively 1982), similar to what was used in the initial SPARC design (Creely *et al.* 2020) but with updated, more conservative, assumptions inspired by higher-fidelity modelling and more recent analysis of similar operational points.

The POPCON for the ARC design point is shown in figure 1. This figure plots the projected operational contours in both volume-averaged density ($\langle n_e \rangle$) and volume-averaged temperature ($\langle T_e \rangle$) space given a set of assumptions about the energy

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

confinement and the assumed profile shapes to calculate the resulting performance. The operational window for ARC based on this analysis is represented by the blue shaded region in figure 1. As stated above, the assumptions made in the ARC POPCON modelling are conservative, using an energy confinement time relative to the $\tau_{98,y2}$ scaling ($H_{98,y2}$) (Doyle *et al.* 1999) of only 0.9. This is especially notable given that the original ARC concept paper and other similar power plant design studies assume H factors greatly exceeding 1.0 (Sorbom *et al.* 2015; Buttery *et al.* 2021). The profile shapes assumed use a set of functional forms meant to approximate realistic profile shapes. Values of a/L_T and a/L_n are assumed to be approximately constant from $r/a = 0.35$ to $r/a = 0.9$. Inside of this region, the normalised gradient scale length profile is assumed to linearly transition to a value of zero on axis; outside of this region, the profile value is considered a free parameter to adjust to match the volume-averaged density and the total stored energy consistent with the assumed energy confinement time. This approach yields parametrised, but more realistic, profile shapes than traditional parabolic profile assumptions. Additional assumptions made are a 50/50 mix of deuterium and tritium with $T_i/T_e = 0.9$, $Z_{eff} = 1.5$, $n_i/n_e = 0.85$, and density peaking is predicted to approximately follow Angioni's empirical density peaking scaling (Angioni *et al.* 2009). The empirical density peaking for these conditions is found to be $n_e(\psi = 0.2)/\langle n_e \rangle = 1.52$ for electrons, with a value of 1.42 being assumed for the main ion (D and T) peaking to allow for some conservatism. As demonstrated in figure 1 the plasma is projected to operate well above the L to H threshold, as calculated with Martin scaling (Martin *et al.* 2008) and well below any β limits. The chosen point in the potential operational window is projected to produce a fusion power of 1133 MW (P_{fus}) at a plasma gain (Q) of ~51 at volume-averaged temperatures and densities of $\langle T_e \rangle = 11.4$ keV and $\langle n_e \rangle = 2.44 \times 10^{20}$ m^{-3}. A more complete list of the operational point parameters can be found in table 1. We note that the analysis discussed in this paper focuses on a steady-state H-mode operational point in ARC. Determination of the operational path (i.e. density and heating programming) needed to access this eventual steady-state H-mode requires time-dependent modelling, which is planned for future work. However, initial calculations suggest that 50 MW of ICRH should be sufficient for H-mode access in ARC. The following sections primarily focus on analysis of core performance and transport. For more information on divertor and heat flux considerations, the reader is referred to Eich *et al.* (2026).

3. ARC performance projections using physics-based modelling

While POPCONs provide a way to rapidly evaluate and visualise plasma conditions within a large design space, the development of modern integrated modelling tool sets enables similar scoping studies to be performed with physics-based modelling. In this section, we briefly describe the suite of tools that were employed to perform modelling of the ARC operational point, discuss results from medium-fidelity modelling and the sensitivity of the reported results to uncertain inputs and assumptions.

3.1. *Description of the modelling set-up*

Medium-fidelity modelling of the ARC design point utilised a range of models to accurately simulate the numerous physical phenomena (transport, magnetohydrodynamics, heating, radiation, current diffusion and fast particles) required for the prediction of ARC performance. In this section, we briefly describe

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

the models employed but the reader is referred to the included references for additional information on each of the models and their implementation.

3.1.1. Coupled TRANSP–PORTALS

TRANSP[1] is a well-known framework for interpretative modelling utilised extensively throughout worldwide fusion community. TRANSP solves the time-dependent, coupled, energy and particle conservation equations while providing models for known sources and sinks that are present in the ARC plasma. Given an initial geometry and a set of profiles, TRANSP performs a fixed boundary equilibrium calculation while evolving current diffusion and modelling a wide range of phenomena. ARC inductive operation at high current results in a large $q = 1$ surface in the ARC design point. Sawteeth were modelled with the Porcelli model (Porcelli, Boucher & Rosenbluth 1996), resulting in a relaxed q profile to just below 1 near the magnetic axis and an inversion radius around mid-radius (see figure 5). ARC will utilise ICRH with a baseline assumption (that may be optimised in the future) of 3 % hydrogen minority heating scheme to provide the entirety of the external heating. The ICRH was modelled via the full-wave TORIC code (Brambilla 1999) and the ohmic heating profile was calculated within TRANSP. The Monte Carlo fast-ion code NUBEAM (Pankin *et al.* 2004) was used for the calculation of fusion alphas.

Kinetic profiles (n_e, T_e and T_i) were predicted using the PORTALS framework (Rodriguez-Fernandez *et al.* 2022; Rodriguez-Fernandez *et al.* 2024*a*) with TGLF SAT2 (Staebler *et al.* 2016) and NEO (Belli & Candy 2008) calculating the turbulent and neoclassical transport, respectively, and a neural network trained specifically for ARC (EPED neural network (EPED-NN), described below) used to provide an edge boundary condition. Following an initial interpretive TRANSP simulation, profile predictions are performed using the sources and sinks from TRANSP while self-consistently calculating radiation, alpha power and exchange during the profile convergence process. The results from these profile predictions are then passed back to TRANSP where it is run interpretively to model the sources (ICRH, ohmic power, alpha heating, etc.), current diffusion and sawteeth. This results in kinetically constrained, magnetic equilibria and source/sink profiles. Results are passed back to PORTALS, the pedestal is updated consistent with changes in the geometry and profiles and the process continues. In practice, 3–5 total iterations of (TRANSP $\rightarrow$ EPED-NN $\rightarrow$ PORTALS) are required before convergence of the equilibrium + transport is complete. This process is similar to that outlined in the STEP framework (Meneghini *et al.* 2020) and provides a relatively rapid ($\sim$4 hours per condition on 32 CPU cores) way to obtain time-independent evaluation of operational points. The results of the coupled TRANSP–PORTALS workflow serve as the basis for all other modelling presented here.

3.1.2. ASTRA

The integrated modelling suite ASTRA (Pereverzev *et al.* 1991) was also used to assess the performance of the nominal ARC operational point and provide a comparison point with other established workflows. Similar to TRANSP, ASTRA solves time-dependent transport equations (one-dimensional transport while considering two-dimensional geometry) while providing models for the plasma sources and sinks. As ASTRA does not natively include the Porcelli sawtooth model or TORIC for evaluation of the ICRH deposition, results from the interpretive TRANSP modelling

[1] https://transp.pppl.gov

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

were utilised in ASTRA. Specifically, the current density profile and heating deposition profiles obtained from TRANSP were fixed in the ASTRA simulations, while transport, pedestal (via EPED-NN) and radiation were evolved self-consistently to obtain the ASTRA results in this paper. For additional details on the ASTRA framework and implementation, the reader is referred to Pereverzev *et al.* (1991).

3.1.3. TORAX

TORAX (Citrin *et al.* 2024) is a JAX-Python implementation of a one-dimensional time-dependent tokamak plasma transport solver with considerations for two-dimensional magnetic geometry, similar to TRANSP and ASTRA. The auto-differentiation capabilities of JAX allow for implicit time stepping schemes for the time-dependent solution, significantly increasing the dynamically chosen time-step sizes from explicit solvers while respecting the plasma dynamics and thus accelerating the simulation. However, this capability demands each of its components to be in a differentiable form (e.g. analytic expressions or neural network surrogates) to fully leverage this accelerated numerical scheme. The TORAX simulations presented here derive their initial and boundary conditions from the converged reference PORTALS–TRANSP simulation, then evolve the n_e, T_e, T_i profiles using QLK-NN (van de Plassche *et al.* 2026) as the turbulent transport model until a new steady state is achieved. Fixed ICRH profiles from the reference TRANSP simulation and internally self-consistent analytic ohmic, fusion, radiation and exchange source calculations provide the balancing terms for the transport equations. Since the core region is expected to be heavily impacted by the presence of sawteeth, and TORAX currently does not implement a sawtooth model, an inner transport patch was applied in these simulations to empirically estimate these effects while allowing convergence towards a stationary solution. This method has been demonstrated and discussed in previous integrated modelling studies with other models (Ho *et al.* 2019; Fajardo *et al.* 2024).

3.1.4. ARC EPED-NN and pedestal assumptions

To provide a boundary condition for the core ($r/a \sim 0.0$ to 0.95) modelling workflows, an EPED-NN was developed around the ARC parameter space. EPED (Snyder *et al.* 2009, 2011) is a well-known model which predicts the pedestal width and height (pressure or average temperature) via calculated criticality conditions for nearly local kinetic ballooning modes (KBM) and non-local peeling–ballooning modes. It has undergone extensive comparison with experiment, typically demonstrating agreement within $\pm 20\,\%$ of experimentally measured pedestal pressures (Snyder *et al.* 2019; Fenstermacher *et al.* 2025). The EPED-NN used in this paper was trained on approximately 14 000 EPED runs performed with variations of EPED inputs around anticipated ARC operational values (B_T, I_p,n_{ped},κ, δ, β_N, Z_{eff}, ϵ, R, n_{sep}/n_{ped} and T_{sep}). As shown in figures 2(*a*) and 2(*b*), this network is found to reproduce the EPED predictions in both pedestal pressure and width in ARC parameter space to a high accuracy (RMSE $= 2.7\,\%$). However, as demonstrated in figure 2(*c*) the neural network is often found to be least accurate at the transition from peeling- to ballooning-limited pedestals. The EPED model accurately predicts the pedestal height and width in Type I edge-localised mode (ELM) and quiescent H-mode (QH) discharges (Snyder *et al.* 2012). However Type I ELM H-modes are potentially an undesirable operational point for ARC due to large bursts of heat and particle flux associated with Type I ELMs, and there is significant uncertainty over whether ARC will have sufficient $E \times B$ shear in the edge for QH operation.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

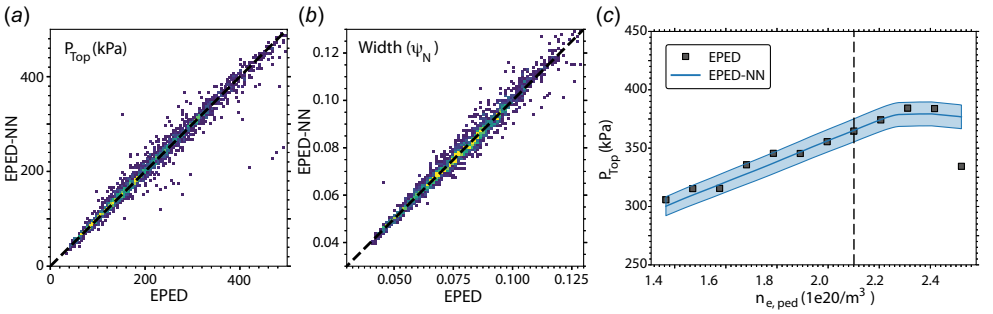


FIGURE 2. Pedestal (*a*) top pressures and (*b*) widths predicted via the ARC EPED-NN are compared with full EPED results. (*c*) A scan of the pedestal density is shown for EPED and EPED-NN with estimated uncertainty in the NN result represented by the shaded region and the vertical line indicating the nominal operational pedestal density.

The EPED model is currently the most validated pedestal model available to the fusion community, and it typically provides an upper limit for the pedestal pressure in regimes other than Type I ELM and QH. For example, ELM suppression with resonant magnetic perturbations is predicted to constrain the width of the pedestal via three-dimensional field penetration near the pedestal top (Snyder *et al.* 2012; Hu *et al.* 2020), resulting in an observed ∼10 %−20 % reduction in the pedestal height from the EPED-predicted value. Similarly, a wide range of improved performance regimes and H-mode varieties have been identified on tokamaks worldwide, many of which are inherently ELM-free and/or offer other potentially favourable pedestal structures (such as I-mode, QCE, EDA H-mode etc.) (Groebner & Saarelma 2023; Fenstermacher *et al.* 2025 and references therein). While the field does not currently possess a complete predictive capability for these regimes, many of these regimes have been demonstrated experimentally to exhibit pedestal pressures that approach those found in type-I ELM H-mode conditions, suggesting that pedestal pressure approaching that predicted by EPED could possibly be attained without large ELMs in future ARC conditions. Definitive confirmation of such conditions will benefit from pedestal model validation in SPARC. We also note that EPED under certain conditions predicts a bifurcation of the pedestal solution, and the existence of a very high 'Super H' pedestal solution at significantly higher pedestal pressure than the usual H-mode solution (Snyder *et al.* 2015). Very high pedestals consistent with Super-H predictions have been observed on Alcator C-Mod and DIII-D (Hughes *et al.* 2018; Snyder *et al.* 2019). In the interest of making conservative projections, Super H solutions are not considered here, but may be explored in future studies. For additional discussion on access to ELM-free regimes, the reader is referred to Eich *et al.* (2026).

Development of new, validated and predictive pedestal models is an open area of research and will be incorporated into future work. For example, the EPED model has recently been extended and coupled to SOLPS to enable coupled pedestal–scrape-off layer (SOL)–divertor predictions, including self-consistent calculation of separatrix conditions and particle sources in the edge plasma, which is a necessary condition for including non-stiff transport models in EPED. Coupled pedestal–boundary calculations are expected to be included in near-future ARC physics studies.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

3.2. *ARC power balance*

Initial analysis of the ARC design point utilised the TRANSP–PORTALS workflow for the calculation of the power balance. The resulting power flows, heating profiles and radiation were calculated self-consistently and provide some insight into the physics of the operational point. The ion and electron power flows calculated via TRANSP–PORTALS calculations are plotted in figures 3(*a*) and 3(*b*) with each contribution to the total flows indicated in the legend. As expected in a power plant class fusion device operating with deuterium and tritium, the dominant source of power in the system is provided via fusion alphas with a majority of fusion power transferred to the electrons (130 MW). Operation with high absolute density implies that collisional exchange also plays an important role in the overall power balance. The simulated profiles in this condition have $T_e > T_i$ leading to an overall power flow from the electrons to the ion channel due to collisional energy exchange, accounting for about 45 MW of total power loss from the electron channel and representing a source of power to the ions that is approximately equal to that coming from fusion products. The applied auxiliary heating is ohmic + ICRH ~ 22 MW. The ICRH contributions are calculated using the TORIC module and it is found that this mechanism provides slightly more ion heating than electron heating (after considering the minority slowing down) in this condition. Since the majority of the power input into the system is into the electron species (via ICRH and the slowing down of fusion products), and the presence of electron temperatures that exceed ion temperatures, it might be expected that total power flows are dominated by electron losses. Perhaps counterintuitively, overall transported power is dominated by the ion channel as shown in figure 3(*d*).

In figure 3 it is found that the ion heat flux exceeds the electron heat flux across the radius, leading to ion-to-electron-heat-flux ratios (Q_i/Q_e) in the 1.5–3 range across the profile. The ratio of power flows in these conditions has important implications for the dominant turbulence as is discussed in § 4 in more detail. Despite the strong electron heating via fusion, these heat flux ratios are obtained by two mechanisms: (i) strong collisional exchange transferring power from the electrons to ions ($\sim$31 MW) and (ii) large radiative losses ($\sim$66 MW) from the plasma due to the presence of high-Z impurities (tungsten) and the high operating temperatures. Figure 3(*c*) plots the total radiation and its contributions as calculated within TRANSP–PORTALS. As shown in this figure, bremsstrahlung, line radiation and synchrotron radiation all play potentially important roles in setting the overall radiation losses.

In figure 3(*a*), one can see the overall impact of radiation on electron power flows. Radiation contributions are large, approaching the power flow into the electrons provided by fusion in overall magnitude. As a result, the overall power flowing through the electron channel is substantially smaller than that of the ions, which do not experience similar losses. It is important to note that this radiation is dominated by the presence of a small tungsten concentration, assumed to be $n_W = 1.5 \times 10^{-5} \times n_e$. This value was chosen to be consistent with previous modelling of devices such as both SPARC (Rodriguez-Fernandez *et al.* 2022) and ITER (Howard *et al.* 2024*b*) and has been shown to be attainable in current tungsten-walled fusion devices (Neu *et al.* 2005, 2016). We note that the overall radiation produced by tungsten is highly dependent on both the assumed concentration and on the tungsten cooling factors (L_z) which can exhibit high uncertainty (Putterich *et al.* 2010). However, these uncertainties are unlikely to be large enough to qualitatively change the ion dominance of the power balance in these conditions.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

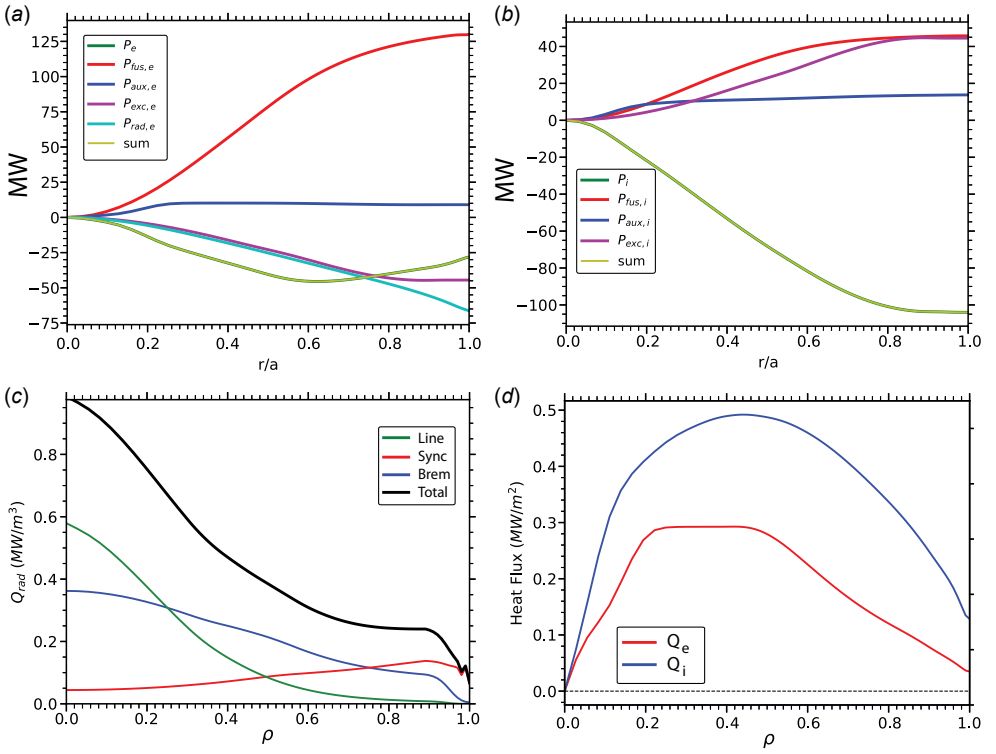


FIGURE 3. The electron power flows (*a*), ion power flows (*b*), radiation contributions (*c*) and heat fluxes (*d*) are plotted for the ARC design point. Power sinks in (*a*,*b*) are given, for visualisation purposes, a negative sign.

3.3. *Projected ARC performance*

Projections of ARC performance were first performed using the TRANSP–PORTALS framework, with results from this analysis then serving as the starting point for analysis with ASTRA, TORAX and PORTALS–CGYRO. All medium-fidelity integrated modelling workflows attempted to utilise similar physics, although a proper 'apples to apples' comparison is not possible given the differences in the codes. Here we attempt to summarise the numerical set-up for each framework. All codes simulated conditions with five assumed species (D, T, H (3 %), lumped impurity and W) with $Z_{eff} = 1.5$. A tungsten concentration of $1.5 \times 10^{-5} \times n_e$ and a lumped impurity of $Z = 5$, $A = 10$ were utilised to provide additional dilution and resulting in $n_{D+T}/n_e = 0.85$. Although seeded impurities will likely be necessary in ARC divertor, the lumped impurity assumed here is not meant to describe any specific impurity but instead to represent some combination of He ash, seeded impurities and any other contaminants. This impurity's Z is chosen to match the assumed dilution (0.85) and Z_{eff}. Self-consistent modelling of seeded impurities and their transport is part of ongoing work. The ICRH power deposition profiles, ohmic heating profiles and current density profiles were obtained from the TRANSP–PORTALS calculations and were fixed during analysis with ASTRA and TORAX. Radiation was calculated consistent with the predicted kinetic profiles but it should be noted that differences in the impurity cooling factors likely contribute to some differences in the total radiation. For transport calculations, ASTRA and

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

TRANSP–PORTALS used TGLF SAT2 for calculation of turbulent driven fluxes (Staebler *et al.* 2016) and TORAX used a neural network derived from the quasi-linear transport model QualiKiz (QLKNN) (van de Plassche *et al.* 2026) which captures only electrostatic turbulence. The TGLF SAT2 setting used for both the TRANSP and ASTRA workflows included electromagnetic turbulence and added resolution at long wavelengths.[2] Neoclassical fluxes, generally found to be negligible, calculated via analytic formula for ASTRA, are calculated with NEO (Belli & Candy 2008) in the TRANSP–PORTALS framework, and are currently not implemented in TORAX. All modelling approaches simulate from axis to approximately $r/a = 0.9$ but there are differences in the core modelling assumptions. TRANSP–PORTALS predictions utilise a relatively sparse radial grid spanning from $r/a = 0.35$ to 0.9, choosing to linearly interpolate gradient-scale lengths to zero inside of $r/a = 0.35$. In contrast, ASTRA and TORAX simulate all the way to the plasma axis, with TORAX applying a core flattening of profiles to approximate the impact of sawteeth. Pedestal pressure and width were obtained from EPED-NN and were self-consistently evolved during the ASTRA and TRANSP–PORTALS simulations to adjust for changes in the calculated β_N (a core parameter that affects the pedestal predictions, thus generating a feedback loop necessary to perform self-consistent predictions) while TORAX utilised fixed pedestal top conditions obtained from the converged TRANSP–PORTALS simulation, as TORAX currently lacks implementation of EPED-NN. The pedestals for both n_e and T_e were assumed to exhibit a mtanh functional form. The final pedestal for the TRANSP–PORTALS modelling predicted a pedestal top pressure of approximately 360 kPa with a pedestal top width of $\Psi_{pol} = 0.074$, which serves as the boundary condition for the core modelling.

As demonstrated in figure 2 the ARC design point sits in the peeling branch of H-mode parameter space (as predicted by the EPED model), about 10 % below the transition from peeling- to ballooning-limited pedestal conditions. The high density of the operational point was chosen to take advantage of the favourable scaling of fusion power with density ($P_{fus} \propto n^2$) while maintaining a safe distance from the peeling–ballooning transition. The input pedestal density (2.1×10^{20} m^{-3}) to EPED-NN was chosen also to a obtain volume-averaged density as close as possible to that specified in the POPCON scoping of $\langle n_e \rangle \sim 2.4 \times 10^{20}$ m^{-3}. However, in 1.5-dimensional modelling, the volume-averaged density is not an input but results from a combination of the assumed pedestal density and profile shape, combined with the density profile predictions in the core. Although not explicitly utilised in core modelling, a ratio of separatrix density to pedestal density ($n_{sep}/n_{ped} = 0.4$) was assumed for the base ARC simulation with a separatrix temperature of 200 eV. These values are inputs to EPED and can play an important role in setting the predicted pedestal pressure, which in turn translates into differences in the core performance. A discussion of the details of these assumptions is given in a paper that is also part of this ARC physics basis collection (Eich *et al.* 2026). Changes in the assumed separatrix density will be investigated as part of future work focused on core-edge integration in ARC conditions.

The T_e, T_i and n_e profiles predicted by TRANSP–PORTALS, ASTRA and TORAX are plotted in figure 4 and the reader is referred again to table 1 for additional performance metrics. Generically, it can be seen that these approaches all yield qualitatively similar profiles, despite the known differences in the modelling approaches and underlying transport models. However, there are clear quantitative

[2]Numerical set-up referred to as SAT2-ASTRA utilises NBASIS_MAX = 6 and KY_GRID = 4.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

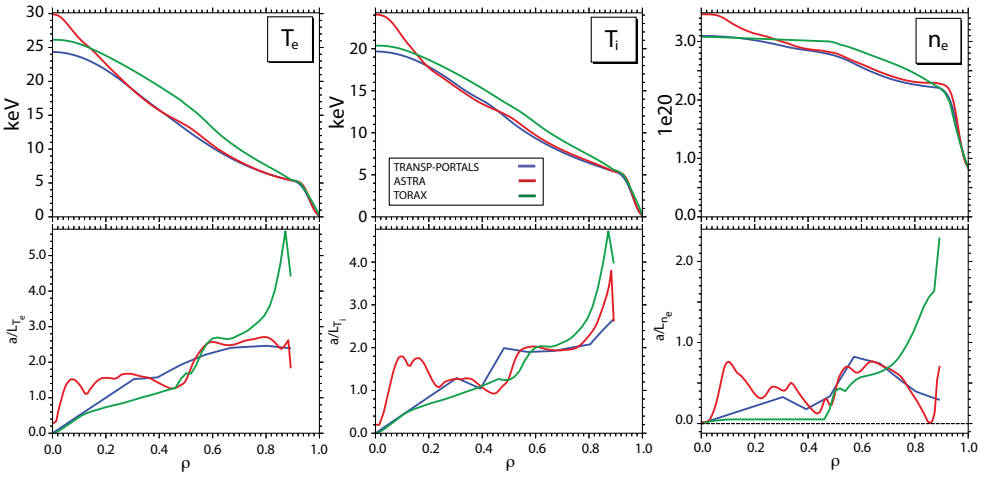


FIGURE 4. The T_e, T_i and n_e profiles and their respective normalised gradient-scale lengths are plotted for predictions from TRANSP–PORTALS (blue), ASTRA (red) and TORAX (green).

differences in the predicted profile shapes, particularly in the predicted electron density. At FPP-relevant temperatures, the highly nonlinear dependence of the fusion power on local values of density and temperature implies that even small changes in the profiles can in principle lead to non-negligible performance changes. Profiles predicted via TRANSP–PORTALS and ASTRA lie in very close agreement, predicting 879 MW ($Q = 38.6$) and 895 MW ($Q = 37.1$), respectively. The small variance in performance likely results in part due the differences in predicted near-axis gradients, which are only interpolated in the TRANSP–PORTALS workflow. In contrast, the modest difference in predicted profile shapes shown in TORAX results in dramatically different performance predictions. Increased density and temperature peaking in the near edge yields broader profiles and predicted fusion power of 1295 MW ($Q = 51.0$). Differences between the integrated modelling results and the POPCON predictions have their origin in the predicted profile shapes relative to those assumed in the 0-D evaluations. As shown in table 1 all medium-fidelity approaches predict similar values of $H_{98,y2}$, with values ranging from 0.9 to 0.94, in close quantitative agreement with the POPCON assumption ($H_{98,y2} = 0.9$). The source of disagreement is found when comparing the volume-averaged quantities and their associated peaking. The POPCON assumptions for profile peaking are significantly higher than those found the TRANSP–PORTALS and ASTRA modelling reflected in particular in the differences in the peaking factors ν_{ne}, ν_{Te} and ν_{Ti} as well as volume-averaged temperatures (recall that $\langle n_e \rangle$ is essentially chosen to be similar between POPCON and medium-fidelity modelling through the choice of n_{ped}). This general result is consistent with modelling published for the SPARC primary reference discharge (Rodriguez-Fernandez *et al.* 2020) which demonstrated that despite similar 0-D agreement between POPCON and medium-fidelity modelling, the predicted performance was significantly lower due to differences in profile shape. All medium-fidelity modelling efforts predict that ARC will operate well below anticipated beta limits (with $\beta_N = 1.58-1.94$) and will operate above the L to H threshold as predicted via the Martin scaling (Martin *et al.* 2008). The range of predicted fusion performance (879–1295 MW) despite nearly identical H factors is impressive

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

Quantity	Units	POPCON	TRANSP + PORTALS	ASTRA	TORAX	PORTALS–CGYRO
P_{fusion}	MW	1133	879	895	1295	677
Q_{plasma}		51.0	38.6	39.1	51.0	29.9
$H_{98,y2}$		0.90	0.90	0.93	0.94	0.87
τ_e	s	0.93	1.11	1.15	0.98	1.26
P_{SOL}	MW	112.0	132.0	148.4	227.7	86.6
P_{SOL}/P_{LH}		1.16	1.12	1.23	1.65	0.75
$n_e(\Psi = 0.2)/\langle n_e \rangle$		1.52	1.27	1.30	1.25	1.18
$T_i(0)/\langle T_i \rangle$		X	2.20	2.64	2.02	2.27
$\langle n_e \rangle$	10^{20} m^{-3}	2.44	2.40	2.47	2.56	2.34
$\langle T_e \rangle$	keV	11.4	9.9	10.1	11.7	9.13
$\langle T_i \rangle$	keV	10.3	8.9	9.1	10.1	8.4
f_{GW}		0.90	0.88	0.90	0.89	0.86
β_N		1.80	1.58	1.51	1.95	1.29

TABLE 1. Summary of performance from various models of the ARC operational point. The P_{LH} utilised is taken from Martin *et al.* (2008).

and suggests that the projection with energy confinement scalings is insufficient for fusion devices operating at power-plant-relevant plasma conditions. The range of integrated modelling approaches applied would suggest that this ARC design point should expect approximately 1 GW of fusion with a plasma gain of ∼40 and are within ±20 % of the POPCON prediction. Even with the most pessimistic of these medium-fidelity modelling predictions, fusion powers approaching 1 GW should still provide sufficient fusion power to constitute a viable design.

As mentioned above, ARC will employ demountable field coils as part of its design, allowing for significant flexibility in device operation, maintenance and even the ability to interchange physical components such as the tokamak vacuum vessel. Additional potential operational points have been evaluated which are possible within the device design. With trade-off between the longevity of the tokamak components and achieved fusion performance, operational points up to 1100 MW have been identified using the TRANSP–PORTALS framework with the nominal physical assumptions and further investigation may be the subject of future work.

3.4. *Design point sensitivity*

To investigate the sensitivity of the medium-fidelity modelling results to input assumptions, a series of TRANSP–PORTALS scans were performed around the ARC reference design point. The dependence of performance in H-mode tokamak discharges on pedestal pressure (predicted fusion power typically scales approximately with the square of the pedestal pressure) is well known in the literature (e.g. Kinsey *et al.* 2011) and all modelling described in previous sections assumed a pedestal width and pressure consistent with the EPED model. For the purposes of this effort, we chose to scan the assumed edge pressure around the nominal EPED prediction – a range that spans the published level (±20 %) of agreement of EPED with experiment (Snyder *et al.* 2019; Fenstermacher *et al.* 2025). This exercise provides insight into the range of pedestals that may be accessible for the

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

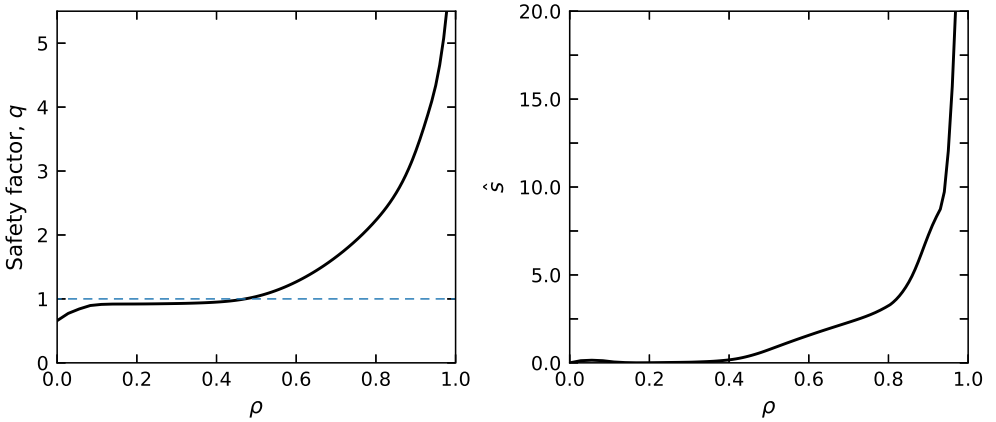


FIGURE 5. The safety factor profile and magnetic shear profile $(-(r/q)\mathrm{d}q/\mathrm{d}r)$ from the TRANSP–PORTALS and CGYRO modelling are shown.

ARC tokamak, including those at potentially degraded pedestal performance – which may be necessary to avoid ELMs during ARC operation, in the absence of sufficiently detached operational conditions. The results of the pedestal top pressure scan are plotted in figure 2(*c*). Unsurprisingly, the results are extremely sensitive to changes in pedestal pressure. A 20 % increase in the pedestal pressure raises fusion performance from 878 MW to ∼1.24 GW ($Q = 55.3$) of fusion power production while raising the predicted $H_{98,y2}$ to 0.98, bringing the results into better agreement with the most commonly used H-mode empirical scaling. However, the need to avoid ELMs in power-plant-class devices is well established, and therefore reduction of the pedestal pressure below the EPED-predicted operational point may be necessary for safe operation of the device. A 20 % decrease in the assumed pedestal pressure drops the predicted performance to approximately 565 MW ($H_{98,y2} = 0.8$), a significant drop in fusion power, but still producing over half a gigawatt via fusion. It is notable that with this drop in performance the predicted $H_{98,y2}$ drops to just 0.8, outside of the quoted $H_{98,y2}$ uncertainty, and the total power through the SOL is likely insufficient to maintain robust H-mode operation. The assumed pedestal top pressure is not the only input assumption that can play a crucial role in setting the overall performance. As outlined in a recent SPARC publication (Muraca *et al.* 2025), assumptions on parameters such as n_W/n_e and $(T_i/T_e)_{ped}$ can also play a significant role. In this work, we have examined the sensitivity of the ARC results to assumed values of n_{sep}/n_{ped}, n_W/n_e and $(T_i/T_e)_{ped}$, and the dependence of results on P_{ped}. These results are plotted in figure 6(*b–d*). As shown in figure 6(*c*), when holding the pedestal density constant and scanning n_{sep}, there is an apparent weak dependence on the assumed n_{sep}/n_{ped} around the reference conditions. However, this is somewhat misleading, as conditions with $n_{sep}/n_{ped} \geqslant 0.45$ actually operate on the ballooning branch compared with peeling-limited operation for $n_{sep}/n_{ped} \leqslant 0.45$ for the nominal value of Greenwald fraction assumed in the operational point ($f_{GW} = 0.9$). This is important because operation on the ballooning branch means that slight increases in density will not increase fusion but will lead to a significant drop in fusion performance. In figure 6(*d*) the dependence of the performance on assumed tungsten concentration over the limited range of $(0.75{-}4.5) \times 10^{-5} \times n_e$ is shown. To remind the reader, the nominal operational point for the ARC discharge is assumed to be

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

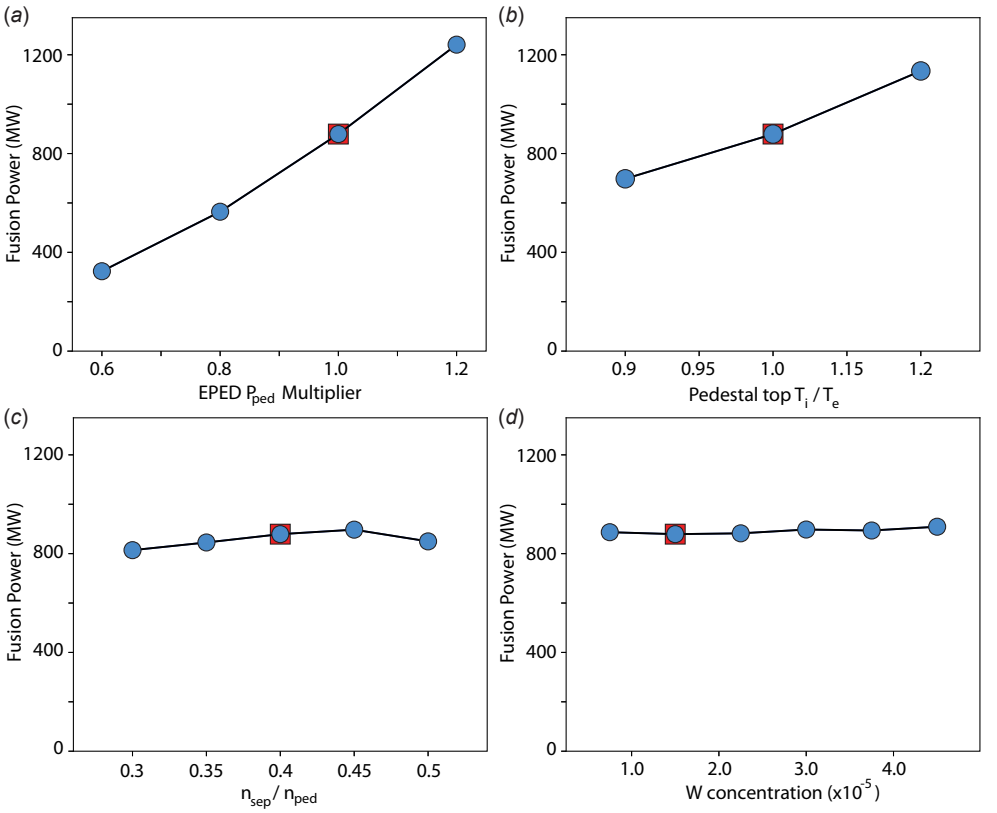


FIGURE 6. (*a–d*) The sensitivity of the TRANSP–PORTALS modelling results to scans in uncertain input parameters is plotted. The reference case is indicated in red.

$1.5 \times 10^{-5} \times n_e$ with recent publications (Ivanova-Stanik, Chmielewski & Zagorski 2023) suggesting that lower concentrations may be obtained due to high-density operation. Despite radiation varying from ∼57 to 110 MW over the range of scanned values, the dependence on this parameter is weak, with variations in performance likely arising from small differences in numerical convergence. This result is perhaps not completely unexpected as similar burning plasmas have been demonstrated to exhibit extremely stiff transport (Rodriguez-Fernandez *et al.* 2020; Holland *et al.* 2023) implying that changes in the transported power are unlikely to significantly modify predicted profiles. Increasing the fidelity of impurity transport and radiation calculations around ARC design points will be the subject of future work. All modelling results presented made the assumption that $T_e = T_i$ at the pedestal top. Although the high density and relatively cool plasma present at the pedestal top will likely lead to significant exchange and therefore approximately equal ion and electron temperatures, $T_e = T_i$, it is unclear to what extent this assumption will hold and the temperature ratio T_i/T_e is known to play an important role in stabilising (or destabilising) ion temperature gradient (ITG) turbulence (Casati *et al.* 2008). As shown in figure 6(*b*) the performance is very sensitive to the assumed pedestal top temperature ratio. A range of 698–1134 MW is found for a variation of T_i/T_e from 0.9 to 1.1. This provides an indication that the dominant turbulence setting ARC core performance is likely ITG – a result that is confirmed in § 4 of this paper.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

Taken together these results suggest that a wide range of performance predictions are possible within modelling uncertainties and underscore the need for an enhanced understanding of pedestal performance and structure to refine estimates for burning plasmas (such as SPARC) and next-generation fusion devices aimed at higher fusion power generation.

4. ARC performance projections using nonlinear gyrokinetics

The analysis described in the previous section relies primarily on what are referred to as physics-based models. Such models are routinely applied to both predict and interpret existing fusion experiments due primarily to their balance of accuracy and computational cost. For this reason, they are commonly used in scoping studies for next-generation fusion devices (Fable *et al.* 2019; Buttery *et al.* 2021). However, recent advances in high-performance computing hardware and the optimisation of high-fidelity turbulence modelling for GPU-based high-performance computing systems have enabled the more routine use of nonlinear gyrokinetics for profile predictions and projections of performance. At the time of writing, nonlinear gyrokinetics is the gold standard for profile predictions, as it is generally believed to be capable of capturing the relevant physical phenomena needed to predict profiles and performance in the core ($r/a < 0.9$) of tokamak plasmas. Given the need for reliable, high-accuracy modelling, this section covers the prediction of ARC performance via nonlinear gyrokinetic profile prediction.

4.1. *Simulation and modelling set-up*

Predictions of plasma profiles utilising nonlinear gyrokinetics were enabled by surrogate-accelerated profile prediction techniques implemented in the PORTALS code (Rodriguez-Fernandez *et al.* 2024*a*). The details of this framework are out of the scope of this paper but the reader is referred to both Rodriguez-Fernandez *et al.* (2022) and Rodriguez-Fernandez *et al.* (2024*a*). However, we summarise the numerical techniques in PORTALS here for completeness. To accelerate the convergence in profile prediction, PORTALS first utilises a small database ($\sim$5) of high-fidelity simulations (nonlinear gyrokinetics) to build a surrogate that is capable of reproducing the response of heat and particle fluxes to inputs (a/L_{T_i}, a/L_{T_e}, a/L_{n_e}, etc.) The surrogates for the nonlinear gyrokinetic fluxes are then used to predict a set of profiles and gradients that are believed to match the transport targets (Q_e, Q_i, Γ_e) when evaluated by the high-fidelity model (nonlinear gyrokinetics). Surrogate-predicted profiles are then tested with the high-fidelity model and the fluxes are reported back to the PORTALS framework. If the fluxes produced are in agreement with the targets, the process concludes and the simulation is deemed converged. However, if the fluxes disagree with the targets, then the results of the high-fidelity model are added to a database, a new surrogate is built and the process continues until convergence is reached between the target fluxes and the results from the high-fidelity model. The surrogate-based approach has been demonstrated to significantly speed up profile predictions compared with traditional methods based on Newton solvers (Rodriguez-Fernandez *et al.* 2024*a*) and has recently been utilised to model experiments on both the DIII-D (Howard *et al.* 2024*b*) and ASDEX-Upgrade (Bielajew *et al.* 2025) tokamaks, demonstrating good agreement with measured profiles and measured turbulent fluctuations.

The high-fidelity evaluations presented in this paper make use of the NEO (Belli & Candy 2008) and CGYRO (Candy, Belli & Bravenec 2016) codes for calculation

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

of the neoclassical and turbulent contributions to the total fluxes, respectively. The total flux in this context is the simple linear sum of the neoclassical and turbulent contributions (i.e $Q_{neo} + Q_{turb}$) with the latter playing the dominant role in all calculations performed for the ARC tokamak. CGYRO is a flux-tube gyrokinetic code that has been optimised for multi-scale turbulence simulation and GPU architectures and has been compared with experiments worldwide (Odstrčil *et al.* 2020; Howard *et al.* 2024*b*; Bielajew *et al.* 2025). Profiles of n_e, T_e and T_i were predicted using the PORTALS framework at five radial locations, $r/a = 0.4$, 0.55, 0.75, 0.875. 0.9, corresponding to values of normalised square root of the toroidal flux, $\rho \sim 0.35$, 0.48, 0.67, 0.8, 0.83. Outside of $r/a = 0.9$, the profiles are held fixed and consistent with the TRANSP–PORTALS modelling described in § 3. The choice of radial evaluation points and their locations is largely based on previous work, which indicates that this choice of radial locations can produce accurate simulation of inductive tokamak plasmas (Rodriguez-Fernandez *et al.* 2024*a*). We note that the innermost location used in this work is at a slightly larger minor radius than previous works (Howard *et al.* 2024*a*,*b*). The extremely low values of magnetic shear and low overall values of q (< 1) present at $r/a = 0.35$ (see figure 5) made these simulations extremely computationally challenging, and thus we chose to move outwards slightly to reduce the computational burden and avoid regions that will be strongly dominated by sawteeth. Evaluation of near-axis locations may be the subject of future work focused on evaluation of fast-ion interactions with turbulence and impurity peaking in near-axis conditions. Unless otherwise noted, all simulations performed in this work were nonlinear, featuring the use of four gyrokinetic species: deuterium, tritium, a lumped impurity ($Z = 4$) and electrons. The validity of the assumption of the lumped impurity has been verified in previous works (Howard *et al.* 2024*a*) and is meant to represent the average Z of several impurities such as the H minority (for ICRH), He ash and W introduced from the metal walls. Despite bundling within the nonlinear gyrokinetic code, the PORTALS framework separates the impurity species when calculating radiation contributions to determine the transport targets. Simulations were fully electromagnetic, evolving $\delta\phi$, $\delta A_{||}$ and $\delta B_{||}$, and included geometry described by the Miller extended harmonic model that is capable of realistic shaping (Arbon, Candy & Belli 2020) and Sugama collisions (Sugama, Watanabe & Nunami 2009). Typical simulation boxes used in this work were approximately $[L_x, L_y] \sim [120, 120]\rho_s$ in the radial and binormal directions and were represented using 512 radial modes (n_r), 24 toroidal modes (up to $k_\theta \rho_s = 1.2$), 24 points in theta (n_θ), 24 pitch angles (n_η) and 8 energies (n_{energy}). These choices of resolutions and simulated range of $k_\theta \rho_s$ are capable of capturing KBM, microtearing mode (MTM), ITG mode and trapped electron mode (TEM) contributions. High-k TEM/electron temperature gradient (ETG) contributions are not captured but are not believed to play an important role as discussed in the following section. Simulation boxes at $r/a = 0.4$ were larger than at other radial locations due to the nature of the local turbulence and low magnetic shear. At this radial location 768 radial modes and 32 toroidal modes were used to simulate box sizes $[L_x, L_y] \sim [400, 160]\rho_s$. Due to the extreme uncertainty associated with prediction of intrinsic rotation profiles, the rotation in all simulations was set to 0. Since it is generally found that finite rotation has favourable effects on the stabilisation of turbulence, this should be viewed as a conservative assumption – but it also may play an important role in future studies of impurity transport and peaking. Recent investigation into the impact of projected intrinsic rotation on SPARC turbulence (Howard *et al.* 2021) found very little impact, but more complete investigations of impact in ARC will be the subject

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

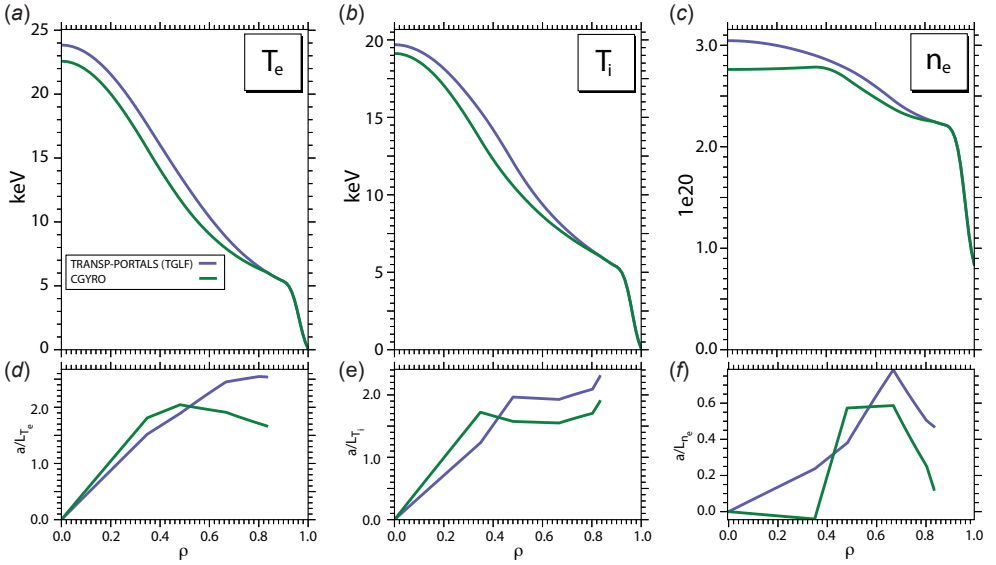


FIGURE 7. The initial TRANSP–PORTALS (TGLF SAT2) profiles (blue) are compared with the converged CGYRO profiles (green) for T_e, T_i, n_e (*a–c*) and the normalised gradient-scale lengths a/L_{T_e}, a/L_{T_i}, a/L_{n_e} (*d–f*).

of future work. Fast ions were not included in this analysis, despite the fact that they will be present with significant populations in a high-gain device such as ARC. Studies focused on fast-ion interactions with turbulence are a cutting-edge area of research (Di Siena *et al.* 2025; Ruiz Ruiz *et al.* 2025) which is too computationally demanding for this publication but will be the subject of future work dedicated to this subject. Each simulated flux was determined from an approximately $500a/c_s$ time average spanning a steady portion of the simulation with typical assessed errors on these fluxes being $1\sigma \sim 10\,\%$.

The simulation set-up described above is capable of capturing low-k turbulence but does not include any impact of contributions from intermediate- or high-k TEM or ETG. The ion-dominated power flows and the associated ion-to-electron-flux ratios (Q_i/Q_e) described in § 3 (figure 3) suggest that this is likely an accurate approximation. Based on the so-called transport 'fingerprints' paradigm proposed by Kotschenreuther and colleagues (Kotschenreuther *et al.* 2019), the heat flux ratios suggest that ITG or KBM will play the dominant role in setting transport in the ARC conditions. As we discuss later in the section, this conclusion appears to be consistent with linear and nonlinear turbulence analysis. However, the use of low-k nonlinear simulations, lack of rotation and fast-ion modelling are approximations which, if relaxed in future work, could lead to differences in the performance projections.

4.2. *Projected ARC performance*

The high-fidelity predictions of the ARC design point are plotted in figure 7. A total of 10 iterations (50 nonlinear gyrokinetic simulations: 10 iterations × 5 radial locations) were required to reach convergence of the steady-state predictions. The resulting T_e, T_i and n_e profiles are plotted in figure 7(*a–c*) with the corresponding gradient-scale lengths for each profile plotted in figure 7(*d–f*). As shown in this

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

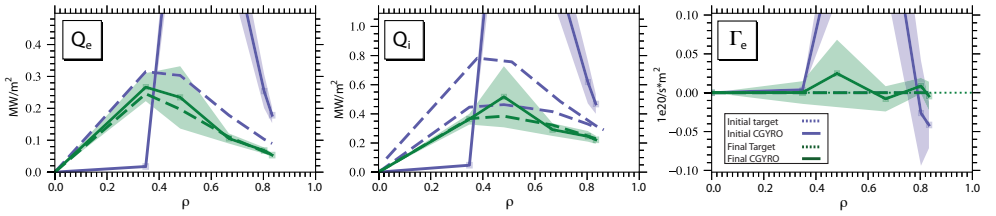


FIGURE 8. The heat and particle fluxes from the initial (blue) and final (green) profiles are plotted with estimated 2σ error bars indicated.

figure the initial profiles (blue) are reasonably close to the final gyrokinetic-predicted profiles (green), suggesting reasonable agreement between predictions with TGLF SAT2 and CGYRO. However, it is important to note that fluxes from the initial profiles (TGLF SAT2) are markedly poorly matched to targets (see figure 8), making the similarity of the final results likely mostly arising from stiff transport. Despite apparent similarity, the changes in profile shapes do lead to significant changes in the overall performance. All profiles (T_e, T_i and n_e) predicted by CGYRO are lower than those predicted by the TGLF modelling with gradient-scale lengths of the ion and electron temperature significantly lower ($\sim$1.5) outside of mid-radius, and only slightly higher at the innermost radial location ($r/a = 0.4$). Consistent with previous work (Howard *et al.* 2024*b*), the density profile demonstrates the largest differences when compared with TGLF SAT2. Smaller values of density gradient-scale length are predicted across the radius. To remind the reader, we note that ARC is source-free (no pellet fuelling or neutral beam injection heating) and therefore the flux-matched density profiles are those required to match a null particle flux condition ($\Gamma_e = 0.0$). Interestingly, the ARC predicted density profile is slightly hollow in the deep core resulting in significantly reduced density peaking and fusion power generation near axis. Recall that the predictions near axis (inside $r/a = 0.4$) are the result of assuming a linear transition of normalised scale lengths to zero on axis and therefore should be viewed as being less certain. However, the small volume in this near-axis region means that its impact on performance is not critical for the overall predictions. As is discussed in the next section, this hollow profile likely results from the presence of unstable KBM in the low-magnetic-shear region of the plasma. Although TGLF SAT2 was generated using results from nonlinear CGYRO simulations, figure 8 clearly shows differences in the predicted fluxes. Such disagreement, however, is perhaps not surprising, as the SAT2 saturation rule was developed from a relatively small dataset of simulations (64) (Staebler *et al.* 2020) and the plasma conditions in ARC and other burning plasmas are generally found to be very stiff. As a result, small differences in predicted gradients can lead to large differences in the predicted fluxes and potentially complicating accurate prediction of profiles.

The predicted lower values of a/L_{T_i} and a/L_n and the resulting T_i and n_e profiles lead to significantly lower performance in these conditions. While predictions from medium-fidelity modelling were typically found to exhibit a wide range of performance of $P_{fus} = [878-1295]$ MW for these conditions, less peaked profiles with smaller volume averages result in a predicted fusion power that falls well short of the POPCON operational point, exhibiting $P_{fus} = 677$ MW with $Q = 29.9$. The drop in the core profiles also results in a significantly lower predicted value of $H_{98,y2}$, energy confinement time and temperature and density peaking factors (see table 1).

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

The calculated $H_{98,y2}$ for this condition is 0.87 making it lower than most data in the database (Verdoolaege *et al.* 2021) while the density peaking factor falls far below the anticipated Angioni value with CGYRO predicting a density peaking $n_e(\Psi = 0.2)/\langle n_e \rangle = 1.18$ compared with an anticipated value of 1.52 obtained from the empirical scaling. The lower density peaking has in recent work been shown to be a consequence of gyro-Bohm transport scaling, the dependence of ITG on collisions and finite-beta effects. It is notable that even the medium-fidelity modelling predicts density peaking values below empirical scalings (table 1), consistent with previous medium-fidelity investigations into projected EU-DEMO conditions (Fable *et al.* 2019). However, the high-fidelity results display even more reduced peaking and thus result in lower performance. Regardless of the modelling performed, this work suggests that ARC (and likely other projected FPPs) should expect to exhibit significantly lower density peaking than empirical scalings. However, despite the lower than desired performance, ARC still exhibits impressive energy confinement, with τ_e of 1.26 s, and operates with a plasma gain (Q_{plasma}) of approximately 30. For additional comparison between the medium- and high-fidelity modelling the reader is referred to table 1. For the CGYRO case, some further optimisation of performance is likely needed to maintain the desired fusion energy and power flow through the SOL consistent with robust H-mode operation. We note that the density profile flattening predicted by CGYRO results in an average density below the design value corresponding to $f_{GW} = 0.9$. Increasing the pedestal density to reach $f_{GW} = 0.9$ will increase the predicted pedestal pressure (see figure 2*c*), and is expected to significantly increase predicted fusion power and Q. Furthermore, the strong sensitivity of predicted performance to pedestal pressure indicates that performance may be recovered from pedestal optimisation via minor modifications to shape, aspect ratio and/or q. A comprehensive optimisation of the coupled core–pedestal system is planned for future work. Sensitivity to pedestal conditions is further studied in the following section.

4.3. *CGYRO sensitivity analysis*

The significant sensitivity that was found of performance results to the pedestal top pressure (figure 6*a*) and the relatively limited ability to accurately predict pedestal performance in new machines motivated additional CGYRO profile predictions. Through modification of the assumed pedestal top temperatures, the assumed pedestal pressure was varied $\pm 20\,\%$ and the profiles of n_e, T_e and T_i were predicted in the same manner as described above. Utilising the existing surrogates allowed for rapid convergence of the profile predictions requiring only nine additional iterations to predict the two new conditions. Consistent with the analysis in § 3, the range of scanned pressure is meant to probe known uncertainties in the EPED model when compared with experiment, and can be viewed as a proxy for assessing the potential benefit of shape and aspect ratio optimisation to increase EPED-predicted pedestal height. The results plotted in figure 9 are largely in line with those from medium-fidelity modelling. Fusion performance ranging from 422 to 960 MW is attainable within these reasonable bounds for the pedestal top pressure with density peaking trending as expected, with higher performance (lower effective collisionality) being slightly more peaked than lower-performing conditions. However, the largest gains are found not in the changes to the predicted density, but instead in the increase in ion temperature resulting largely from raising the boundary conditions with otherwise self-similar profile shapes due to strong ITG transport and associated stiff

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

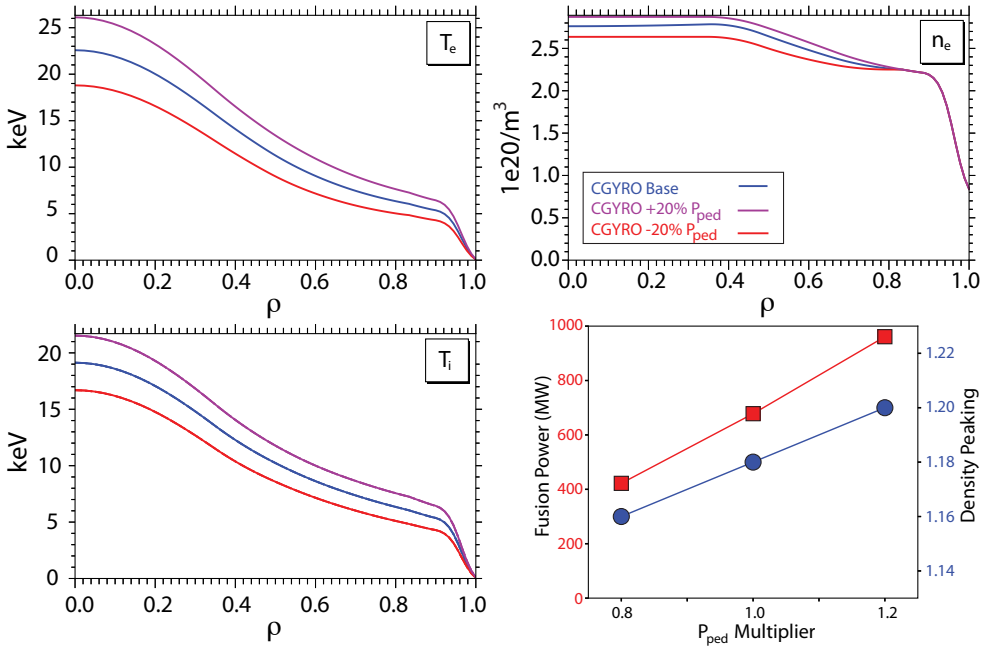


FIGURE 9. The CGYRO-predicted n_e, T_e and T_i profiles obtained for $\pm 20\,\%$ changes in the pedestal pressure are plotted along with the associated fusion power and density peaking.

transport. Interestingly, the 20 % increase in pedestal pressure moves the projected $H_{98,y2}$ from 0.87 (in the reference) to 0.95, improving agreement with empirical scalings and raising the total power flow through the edge to 14 % above the L to H threshold. These results suggest that, even with high-fidelity modelling, a > 500 MW range of fusion performance may be obtained around the ARC operational point, dependent strongly on the observed pedestal performance.

4.4. *Turbulence and transport in ARC*

The profiles predicted via CGYRO + NEO modelling provide insight into the performance that may be achieved in the ARC device. However, additional analysis is required to better understand the nature of the turbulence in this condition. Such investigations are of academic interest, but also are required to understand the necessary physics that must be retained for accurate simulation of ARC and similar tokamak-based FPPs. These results can also provide meaningful targets for investigations on current experiments and SPARC that will enable more relevant validation of existing and emerging transport models. CGYRO was used to perform linear stability analysis at three radial locations, $r/a = 0.4$, 0.75 and 0.9, to provide a look into the linear instabilities and their radial variation in the base ARC conditions predicted by CGYRO. The results of this analysis are plotted in figure 10.

At long wavelengths ($k_\theta \rho_s < 1.0$) turbulent modes rotating in the ion diamagnetic drift direction that are sensitive to changes in a/L_{T_i} are found to dominate at all radial locations. Outside of $r/a = 0.4$, these modes are identified as ITG modes, a conclusion that is largely consistent with the observed heat flux ratios (Q_i/Q_e) and transport 'fingerprints' arguments (Kotschenreuther *et al.* 2019). However, the many low-k modes at $r/a = 0.4$ take on a more electromagnetic character and are

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

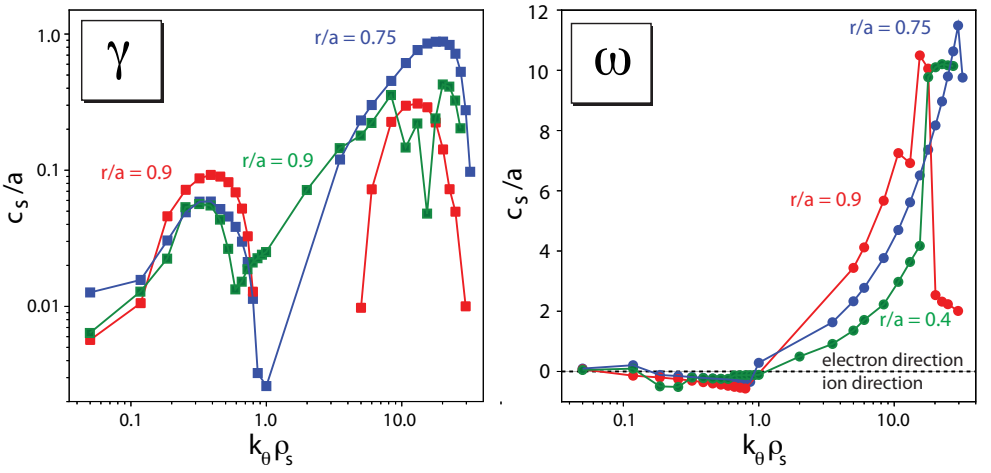


FIGURE 10. The linear growth rates (γ) and real frequency (ω) of the most unstable linear modes at $r/a = 0.4$ (green), 0.75 (blue) and 0.9 (red) are plotted versus $k_\theta \rho_s$ for ion and electron scales.

identified as KBM, destabilised by the low magnetic shear ($\hat{s} = 0.062$), relatively high pressure gradients and moderate values of electron beta ($\beta_e = 0.81\,\%$) at the radial location. This observation is consistent with the non-negligible electron heat flux driven in the nonlinear simulations, where at $r/a = 0.4$ approximately 30 % of the total electron heat flux arises from $A_{||}$ contributions. This is in contrast to all other radial locations, where the heat and particle fluxes are almost completely electrostatic. All radial locations exhibit modes propagating in the electron direction at the largest scales ($k_\theta \rho_s < 0.1$). These modes exhibit tearing parity (odd $\delta\phi$, even $\delta A_{||}$) and are identified as MTM. Despite their presence as unstable linear modes, their overall impact on the saturated turbulence state appears to generally be minimal in the nonlinear simulation at most radial locations, consistent with previous observations from ASDEX-upgrade conditions where both ITG and MTM coexisted (Doerk *et al.* 2015). At $k_\theta \rho_s > 1.0$, intermediate- and high-k modes associated with TEM and ETG turbulence are unstable at all radial locations $r/a = 0.4$, 0.75 and 0.9. However, their presence does not imply the need for simulations spanning both ion and electron scales. Applying existing rules of thumb regarding the importance of high-k TEM/ETG turbulence in the presence of unstable low-k turbulence from both Howard *et al.* (2016) and Creely *et al.* (2019), we conclude that the impact of high k is likely negligible. Both the ratio of high-k to low-k ($\max(\gamma_{high\text{-}k})/\max(\gamma_{low\text{-}k})$ growth rates and the comparison of low- and high-k growth rates when normalised to the $k_\theta \rho_s$ suggest the impact of multi-scale turbulence is likely small and is similar to conclusions for FPP-relevant predictions in recent works (Howard *et al.* 2021; Rodriguez-Fernandez *et al.* 2022; Holland *et al.* 2023; Howard *et al.* 2024*a*).

This analysis provides insight into the linearly unstable modes in ARC but does not provide any knowledge of what exists in the turbulent state which is responsible for heat and particle losses. In this work, we have taken advantage of the unique features of the surrogate modelling to examine the sensitivities of the nonlinear gyrokinetic simulation fluxes (Q_i, Q_e, Γ_e) to changes in turbulence drive terms, a/L_{T_i}, a/L_{T_e} and a/L_n. This analysis essentially amounts to a scan of the relevant drives for the

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

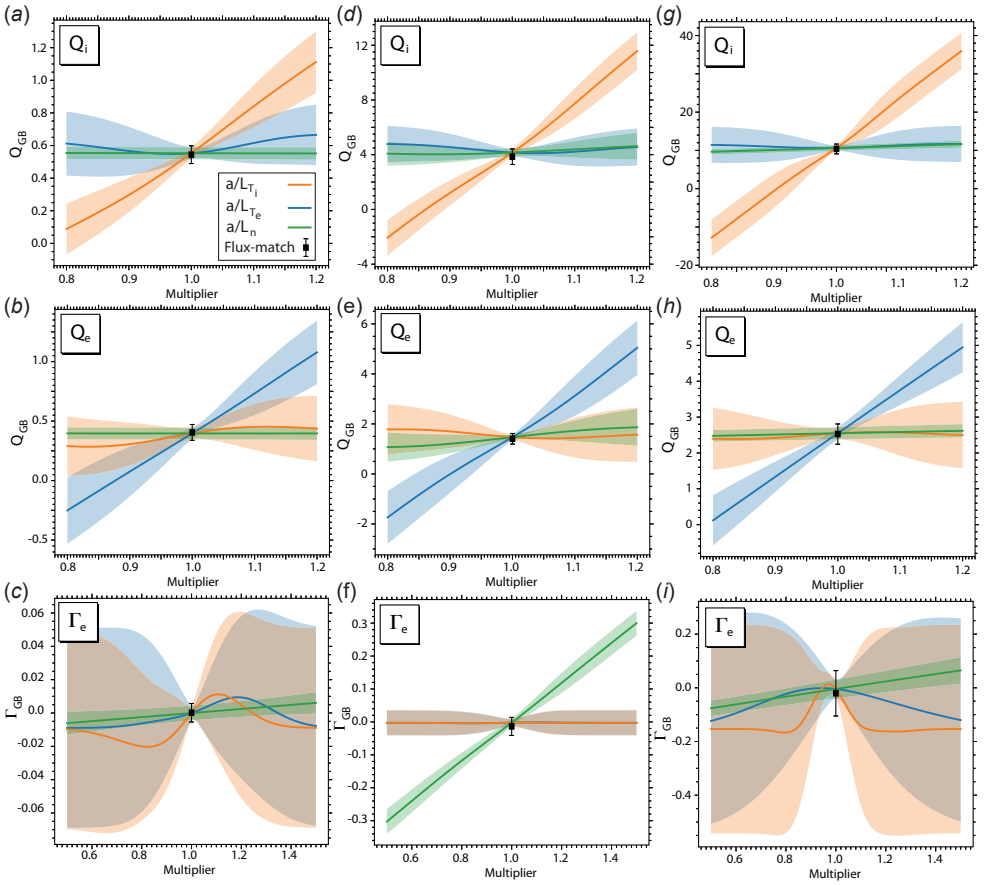


FIGURE 11. The nonlinear gyrokinetic surrogates are plotted around the converged PORTALS–CGYRO simulation for Q_i, Q_e and Γ_e at $r/a = 0.4$ (*a–c*), 0.75 (*d–f*) and 0.9 (*g–i*) for scans of primary drive terms a/L_{T_i} (orange), a/L_{T_e} (blue) and a/L_n (green).

plasma turbulence and can provide more direct insight into the type of turbulence in these conditions. However, instead of using dedicated single-parameter scans, this information in obtained from the nonlinear flux surrogate model and can be rapidly evaluated following a profile prediction. Figure 11 plots the surrogate of the nonlinear gyrokinetic simulations at three radial locations: $r/a = 0.4$, 0.75 and 0.9. It should be noted that most of the profile variations in the surrogate training set exist within a narrow range of the flux-matched condition ($\sim \pm 20\,\%$). Therefore, evaluations of the fluxes outside this range are likely significant extrapolations and can lead to non-physical fluxes (such as negative values for heat fluxes). However, interpreting the results of this analysis to the region with significant training data (indicated by small error bars) yields useful insight into the response of the fluxes to changes in the turbulence drives.

All radial locations simulated indicate that the ion heat flux is most sensitive to changes in a/L_{T_i}, consistent with the linear stability analysis presented in figure 10. This conclusion is consistent with the presence of ITG-dominated turbulence and with recent simulations of ITER and FPP-relevant fusion devices (Rodriguez-Fernandez *et al.* 2022; Holland *et al.* 2023; Howard *et al.* 2024*a*).

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

Interestingly, the electron heat fluxes from the surrogates display a dominant dependence on a/L_{T_e}. Such a dependence would generally be attributed to ∇T-driven TEM turbulence, but this conclusion is inconsistent with the linear stability in figure 10, the nature of turbulence implied by the heat flux ratios in this plasma, and is in striking contrast to recent ITER results (Howard *et al.* 2024*b*). To understand this dependence, a dedicated scan of a/L_{T_e} was performed around the flux-matched conditions (± 10) % which indicates a much weaker dependence on the TEM drive. In fact, the apparent strong dependence of a/L_{T_e} results from the extremely tight coupling of a/L_{T_e} and a/L_{T_i} during the profile prediction convergence. The strong collisional coupling and inter-dependence of the transport lead to a very strong correlation between the variations of a/L_{T_i} and a/L_{T_e} during the training of the surrogate models in these PORTALS simulations. As a result, the surrogate model is most likely unable to distinguish between the dependence on the two parameters particularly given the extremely small training database (10 transport simulations per radial location) that was required to achieve convergence. This is highlighted here to demonstrate a potential weakness of reliance completely on surrogates for interpretation of turbulence characteristics. Particle fluxes (figure 11*c*, *f*, *i*) predicted by the surrogates display a positive dependence on the value of the normalised density gradient-scale length, a/L_n, which is consistent with diffusive particle transport. Somewhat more complicated dependencies on a/L_{T_i} and a/L_{T_e} exist at $r/a = 0.4$ and 0.9. However, the surrogates at these locations are only accurate very close to the flux-matched point, as indicated by the extremely large uncertainties when moving away from this matched point. The conclusion of this work is that the nonlinear results are generally consistent with those found linearly. The ITG turbulence plays a dominant role across most of the plasma profile with some impact of electromagnetic modes (KBM) at the innermost radial locations. These results provide some unique insights into the type of turbulence that should be studied in current and planned devices, to help improve future predictions of ARC operation and highlight that ARC turbulence is similar to that which has been extensively studied and modelled in many existing fusion devices.

5. Conclusions and discussion

This paper outlined the performance and transport predicted for the ARC V3A design point. Modelling using POPCON, performed with conservative assumptions on confinement ($H_{98,y2} = 0.9$) and a more realistic parametrisation (than a parabolic profile assumption) for background profiles, was used to determine the initial design point: a plasma operated with $\langle n_e \rangle = 2.44 \times 10^{20}$ m^{-3} and $\langle T \rangle = 11.4$ keV is projected to produce 1.133 GW of fusion power with an anticipated $\sim$400 MW of electrical production. This design point was studied using a wide range of integrated modelling tools. Initial predictions of the performance from 1.5-dimensional modelling were provided via a new coupled TRANSP–PORTALS framework, with results from this workflow serving as the starting point for additional analysis. Predicted fusion power from the TRANSP–PORTALS workflow was found to be in good agreement with modelling using ASTRA (879 versus 895 MW) with TORAX, predicting significantly higher performance (1295 MW) due to increased density and temperature peaking. These results are in reasonably good agreement with the 0-D POPCON modelling, generally falling within ± 20 % of the POPCON result. However, performance results are extremely sensitive to input assumptions. Values of P_{ped}, T_i/T_e at $(r/a = 0.9)$, n_W/n_e and n_{sep}/n_{ped} were scanned around

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

their nominal assumptions to investigate the sensitivity of the results to known, uncertain inputs. Unsurprisingly, the assumed pedestal top pressure represents a significant lever on fusion performance, producing a range of fusion powers from 565 to 1240 MW of fusion with only a ±20 % variation around the EPED-predicted pedestal pressure. This indicates that a strong performance benefit can be gained by further optimisation of the coupled core–pedestal system, for example via small modifications of the shape and aspect ratio. A similarly large variation (>400 MW) in ARC performance can result from changes of only ±10 % in the assumed pedestal top T_i/T_e ratio, which is a consequence of an ITG-dominated plasma core and stiff transport in high-performance conditions. Analysis was also performed with nonlinear gyrokinetic profile predictions using the PORTALS framework. This work predicted the profiles of T_e, T_i and n_e from $r/a = 0.9$ to the axis utilising a combination of nonlinear CGYRO and NEO simulations. Performance predictions from these high-fidelity investigations fell significantly below the medium-fidelity results. Predictions of profile peaking, in both the temperature and, more dramatically, the electron density, differ substantially from both the medium-fidelity modelling and the POPCON assumptions. These differences lead to much reduced overall fusion power predictions with values of only 677 MW predicted for the nominal operational point. However, it was also shown that with this high-fidelity modelling, even ±20 % differences in the assumed pedestal pressure can lead to 500 MW differences in fusion power (422–960 MW). Furthermore, the predicted flat density profile results in an average density below the design value of $f_{GW} = 0.9$, and raising the pedestal density to recover this design value is expected to significantly recover performance. This result underscores the critical need for accurate pedestal modelling, and the substantial benefit of pedestal optimisation, on predicted performance. Turbulence and transport investigations in these plasmas suggest that the core of ARC is generally dominated by ITG turbulence – but has non-negligible contribution from KBM in the low-shear, high-beta region near the core ($r/a = 0.4$). However, the stability of the KBM is tightly linked to the q profile and therefore different phases of the sawtooth period may exhibit slightly different characteristics. Further investigations will be the subject of future work and may motivate the use of CGYRO (nonlinear gyrokinetic) profile predictions, along with comprehensive pedestal optimisation employing an updated version of EPED, to optimise performance in future ARC design iterations. These results were largely consistent with previous works concerning FPP-relevant plasma conditions and provide a model for the plasma conditions that should be studied in SPARC to improve ARC projections.

5.1. *Sensitivity of results*

The sensitivity of the ARC performance predictions to input assumptions was outlined in the proceeding sections, where it was shown that >500 MW changes in predicted fusion can result from reasonable changes in input assumptions. It is expected that similar sensitivities would be found in most FPPs operating with high Q. The is due to the nonlinear dependence of fusion power on changes in density profiles and temperature profiles. However, additional sensitivities exist that warrant some discussion in the context of performance predictions. Much of the modelling in this paper utilised the physics-based model TGLF SAT2, which has been specifically tuned to reproduce a database of nonlinear CGYRO simulations (Staebler *et al.* 2020). However, additional proposed saturation rules exist (SAT0 (Staebler *et al.* 2007), SAT1 (Staebler *et al.* 2016) and SAT3 (Dudding *et al.* 2022))

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

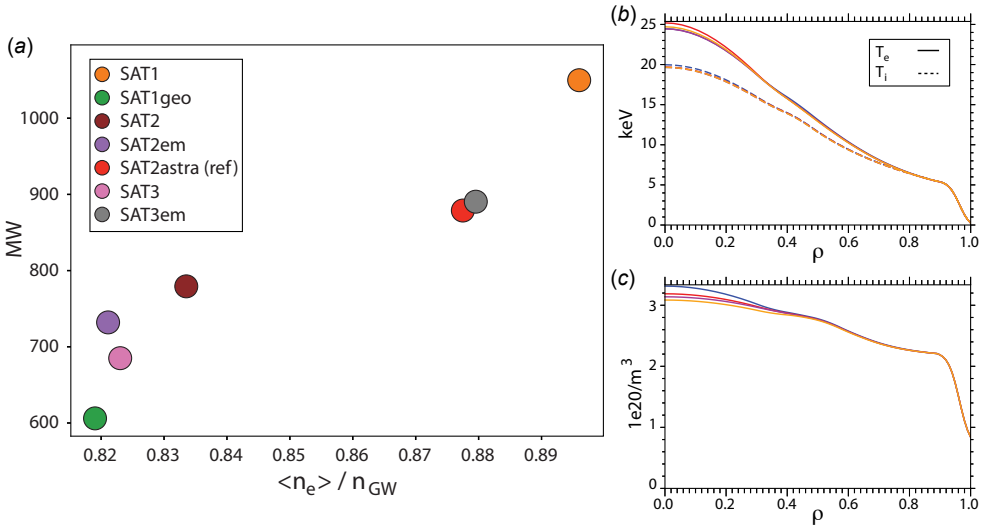


FIGURE 12. The sensitivity of the fusion power predictions to choice of TGLF modelling is shown (*a*). Ion and electron temperature profiles (*b*) and electron density profiles (*c*) are shown for different variations of converged profiles from the TRANSP–PORTALS workflow.

and it is worth noting that the choice of saturation rule, choice of electrostatic versus electromagnetic simulation and even small changes in the numerical set-up of these models can result in non-negligible changes in profiles and performance. This is demonstrated in figure 12 which plots predicted performance for a range of TGLF saturation rules (1, 2 and 3) and numerical settings. SAT0 was omitted from this exercise as it does not include a model for the Dimits shift which is known to dramatically effect the performance in ITG-dominated plasma conditions (Staebler *et al.* 2016). The results of this exercise reveal a 500 MW spread in predicted performance and significant changes in the predicted density profiles as indicated by the nominal $\langle n_e \rangle / n_{GW}$ obtained in the converged profile predictions. Figures 12(*b*) and 12(*c*) plot four different simulations of the identical condition. These simulations are initialised with a random seed, leading to slight changes in the predicted results that all satisfy convergence criteria of the integrated modelling. Despite small changes in the profiles, predicted performance in the four plotted profiles varies by nearly 50 MW. Since fusion power predictions are approximately proportional to $n_i^2 \times T_i^2$, for deuterium–tritium plasmas at fusion-relevant temperatures, predictions of fusion performance are extremely sensitive to predicted profile shapes. This is a feature of the ARC V3A design point that should be anticipated to be present in any realistic tokamak-based FPP design. These results set stringent requirements on modelling to enable accurate performance predictions in tokamak power plants, and also highlight the potential benefits of coupled core–pedestal optimisation.

5.2. *Outlook for SPARC operation*

This paper presents a comprehensive investigation into the performance of the ARC V3A design point. A range of models were applied to predict the performance and transport in ARC plasma conditions, with results pointing to potentially new physics regimes on the horizon. More importantly, these results can provide a

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

compass to guide experimental investigation on current and next-generation fusion devices such as SPARC. This work identifies at least five areas of research that should be pursued on SPARC which are needed to improve predictions for ARC:

(i) Probe the validity of energy confinement time scalings in SPARC.
(ii) Improve understanding and prediction of near-edge and pedestal conditions.
(iii) Investigate density peaking in ARC-relevant plasma cores.
(iv) Validate core transport models in ITG- and KBM-dominated plasmas.
(v) Understand the effects of rotation, impurity transport/peaking and fast ions on performance predictions.

This work and recent burning-plasma-focused publications (Howard *et al.* 2024*a*; Muraca *et al.* 2025) have demonstrated predictions of normalised energy confinement (H factors) well below the empirical scalings. Confinement degradation relative to $\tau_{98,y2}$ and even τ_{89p} (Rodriguez-Fernandez *et al.* 2024*b*) is believed to result, in part, from the extremely stiff nature of transport in the hot-plasma conditions attained in future devices, which has the effect of pinning profiles to the critical gradient of the unstable turbulence. SPARC can provide valuable insight into this potential breakdown of empirical scalings in more burning-plasma-relevant conditions, i.e. those exhibiting long energy confinement times relative to energy exchange times and significant radiation. Such conditions are difficult to access in current devices. SPARC results will help better understand the appropriate assumptions that should be made in future power plant scoping, probe the limitations of existing empirical scalings, contribute to updating of energy confinement scalings and enable validation and improvement of physics models.

As demonstrated in §§ 3 and 4 (figures 6 and 9), near-edge and pedestal conditions represent the most important open question for ARC operation, with reasonable assumptions of pedestal properties able to create a variation of a factor of 2 in predicted fusion performance. The strongest lever to reduce uncertainty in ARC predictions, and further optimise ARC performance, will be the pursuit of systematic studies on SPARC to understand and predict attainable pedestal pressures and T_i/T_e at the pedestal top. The further development of empirical and physics-based models that can accurately predict SPARC pedestals should be pursued and coupled with core-integrated modelling workflows to allow for rapid evaluation of core–edge solutions. Due to power handling considerations not discussed in this paper, it is likely that emphasis on small ELM or ELM-free regimes should be pursued.

Second only to the uncertainty in the pedestal performance, the attainable density peaking represents the largest uncertainty on ARC core performance. Deviations from the empirical scalings have been reported in multiple next-generation devices (Fable *et al.* 2019; Howard *et al.* 2024*a*) and are related in part to the higher attainable beta and ITG dominance of the plasma conditions. Relevant study of this physics requires both low collisionality and ITG-dominated conditions, making it difficult to access in existing devices. SPARC studies will access the relevant turbulence regimes to either confirm or refute the predictions of gyrokinetics and physics-based models, providing a clearer picture of anticipated peaking in ARC and FPP conditions.

Although both physics-based models and nonlinear gyrokinetic models have been validated on devices worldwide, the conditions present in current devices often exhibit properties dissimilar to those predicted in ARC and projected FPPs. Unlike current devices which often exhibit a variety of modes contributing significantly to

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

turbulent transport (MTM, KBM, ITG, TEM, ETG), predictions from next-step devices consistently demonstrate the presence of strong core ITG turbulence (plus potential deep-core KBM) with little to no externally driven rotation, while simultaneously exhibiting very low collisionality. SPARC's operation will provide access to conditions by obtaining well-coupled ions and electrons with long energy confinement times, significant radiation losses and intrinsic rotation expected to be dominated by ITG in most parts of operational space (Rodriguez-Fernandez *et al.* 2022, 2024*b*). Low-q_{95}, H-mode conditions in SPARC will represent the best approximation to anticipated ARC conditions. Extensive validation of physics-based models and nonlinear gyrokinetic approaches should be performed on SPARC to more definitely identify both the models and numerical settings (see figure 12) that accurately reproduce SPARC profiles, allowing for reduced uncertainty in projections of ARC and increased confidence in model-based optimisation of performance.

Lastly, rotation, a detailed study of impurity transport and peaking and fast ions were neglected this work but will likely play a non-negligible role in setting ARC performance. SPARC will enable first-of-kind investigations of intrinsic rotation at high input power (25 MW) and high performance (burning plasma conditions). The peaking of impurities in burning plasmas, including both near-axis peaking and screening in the pedestal region will be crucial to understand on SPARC for projection to ARC. It is anticipated to be the first tokamak burning plasma and the associated high-Q operation represents an ideal testbed for studying interactions of fast ions with turbulence and the impact of fast-ion-driven modes on alpha redistribution and heating.

The results in this paper are the first in-depth look into conditions projected for a high-field tokamak power plant. This work points towards open avenues of research that should be pursued on SPARC and existing devices to reduce performance projection uncertainties, and clarify avenues for further performance optimisation, to help ensure the success of ARC and future FPPs.

Acknowledgements

The authors would like to thank J. Candy and E. Belli for the development of the CGYRO and NEO codes. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the US Department of Energy under contract no. DE-AC02-05CH11231 using NERSC award FES-ERCAP0032045. This work was funded by Commonwealth Fusion Systems under MIT RPP020. This material is based upon work supported by the US Department of Energy, Office of Science, Fusion Energy Sciences, under the Milestone-Based Fusion Development Program under award number DE-SC0024885.

Editor Troy Carter thanks the referees for their advice in evaluating this article.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

specific commercial product, process or service by trade name, trademark, manufacturer or otherwise does not necessarily constitute or imply its endorsement, recommendation or favouring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Declaration of interests

J.C.H., A.C, T.E. and T.B. are employed by Commonwealth Fusion Systems. The authors report no conflict of interest.

REFERENCES

ANGIONI, C. *et al.* 2009 Particle pinch and collisionality in gyrokinetic simulations of tokamak plasma turbulence. *Phys. Plasmas* **16**, 060702.

ANGIONI, C., MCDERMOTT, R.M., FABLE, E., FISCHER, R., PÜTTERICH, T., RYTER, F., TARDINI, G., the ASDEX Upgrade Team 2011 Gyrokinetic modelling of electron and boron density profiles of h-mode plasmas in ASDEX Upgrade. *Nucl. Fusion* **51**, 023006.

ARBON, R., CANDY, J. & BELLI, E.A. 2020 Rapidly-convergent flux-surface shape parameterization. *Plasma Phys. Control. Fusion* **63**, 012001.

BELLI, E.A. & CANDY, J. 2008 Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics. *Plasma Phys. Control. Fusion* **50**, 095010.

BIELAJEW, R. *et al.* 2025 Gyrokinetic profile prediction and validation of a negative triangularity plasma in asdex upgrade. *Nucl. Fusion* **65**, 086042.

BONANOMI, N., MANTICA, P., CITRIN, J., GOERLER, T. & AND, B.T. 2018 Impact of electron-scale turbulence and multi-scale interactions in the JET tokamak. *Nucl. Fusion* **58**, 124003.

BRAMBILLA, M. 1999 Numerical simulation of ion cyclotron waves in tokamak plasmas. *Plasma Phys. Control. Fusion* **41**, 1.

BUTTERY, R. *et al.* 2021 The advanced tokamak path to a compact net electric fusion pilot plant. *Nucl. Fusion* **61**, 046028.

CANDY, J., BELLI, E. & BRAVENEC, R. 2016 A high-accuracy eulerian gyrokinetic solver for collisional plasmas. *J. Comput. Phys.* **324**, 73–93.

CANDY, J. & WALTZ, R.E. 2003 Anomalous transport scaling in the DIII-D tokamak matched by supercomputer simulation. *Phys. Rev. Lett.* **91**, 045001.

CASATI, A. *et al.* 2009 Turbulence in the tore supra tokamak: measurements and validation of nonlinear simulations. *Phys. Rev. Lett.* **102**, 165005.

CASATI, A., BOURDELLE, C., GARBET, X. & IMBEAUX, F. 2008 Temperature ratio dependence of ion temperature gradient and trapped electron mode instability thresholds. *Phys. Plasmas* **15**, 042310.

CASSON, F. *et al.* 2013 Validation of gyrokinetic modelling of light impurity transport including rotation in asdex upgrade. *Nucl. Fusion* **53**, 063026.

CITRIN, J. *et al.* 2014 Ion temperature profile stiffness: non-linear gyrokinetic simulations and comparison with experiment. *Nucl. Fusion* **54**, 023008.

CITRIN, J. *et al.* 2017 Tractable flux-driven temperature, density, and rotation profile evolution with the quasilinear gyrokinetic transport model qualikiz. *Plasma Phys. Control. Fusion* **59**, 124005.

CITRIN, J. *et al.* 2024 Torax: a fast and differentiable tokamak transport simulator in jax. arXiv: 2406.06718.

CREELY, A.J., GREENWALD, M.J. & BALLINGER, S.B. *et al.* 2020 Overview of the sparc tokamak. *J. Plasma Phys.* **86**, 865860502.

CREELY, A.J., RODRIGUEZ-FERNANDEZ, P., CONWAY, G.D., FREETHY, S.J., HOWARD, N.T. & AND, A.E.W. 2019 Criteria for the importance of multi-scale interactions in turbulent transport simulations. *Plasma Phys. Control. Fusion* **61**, 085022.

DI SIENA, A. *et al.* 2025 Understanding turbulence suppression in jet d-t plasma with highly energetic fast ions via global gyrokinetic gene simulations. *Nucl. Fusion* **65**, 086019.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

DOERK, H., DUNNE, M., JENKO, F., RYTER, F., SCHNEIDER, P.A. & WOLFRUM, E. 2015 Electromagnetic effects on turbulent transport in high-performance asdex upgrade discharges. *Phys. Plasmas* **22**, 042503.

DOYLE, E.J. *et al.* 1999 Chapter 2: plasma confinement and transport. *Nucl. Fusion* **39**, 2175.

DUDDING, H., CASSON, F., DICKINSON, D., PATEL, B., ROACH, C., BELLI, E. & STAEBLER, G. 2022 A new quasilinear saturation rule for tokamak turbulence with application to the isotope scaling of transport. *Nucl. Fusion* **62**, 096005.

EICH, T.H., BODY, T.A.J., LOOBY, T.P., BALLINGER, S.B., CREELY, A.J., HILLESHEIM, J.C., SNYDER, P.B., HOWARD, N.T., MASLINE, R. & WIGRAM, M. 2026 Power and particle exhaust for the ARC fusion power plant. *J. Plasma Phys.* **92**, E66. DOI: 10.1017/S002237782610155X

FABLE, E. *et al.* 2019 The role of the source versus the collisionality in predicting a reactor density profile as observed on asdex upgrade discharges. *Nucl. Fusion* **59**, 076042.

FAJARDO, D., ANGIONI, C., KIM, S.H., KOECHL, F., FABLE, E., LOARTE, A., POLEVOI, A., & the ASDEX Upgrade Team 2024 Theory-based integrated modelling of tungsten transport in iter plasmas. *Plasma Phys. Control. Fusion* **67**, 015020.

FENSTERMACHER, M., BAYLOR, L. & DE LA LUNA, E. *et al.* 2025 Progress in pedestal and edge physics: Chapter 3 of the special issue: on the path to tokamak burning plasma operation. *Nucl. Fusion* **65**, 053001.

FIELD, A.R., DUNAI, D., GHIM, Y-C, HILL, P., MCMILLAN, B., ROACH, C.M., SAARELMA, S., SCHEKOCHIHIN, A.A., ZOLETNIK, S. & the MAST Team 2014 Comparison of bes measurements of ion-scale turbulence with direct gyro-kinetic simulations of mast l-mode plasmas. *Plasma Phys. Control. Fusion* **56**, 025012.

GORLER, T., WHITE, A.E., TOLD, D., JENKO, F., HOLLAND, C. & RHODES, T.L. 2014 A flux-matched gyrokinetic analysis of DIII-D l-mode turbulence. *Phys. Plasmas* **21**, 122307.

GROEBNER, R.J. & SAARELMA, S. 2023 Elements of h-mode pedestal structure. *Plasma Phys. Control. Fusion* **65**, 073001.

GUTTENFELDER, W. *et al.* 2013 Progress in simulating turbulent electron thermal transport in nstx. *Nucl. Fusion* **53**, 093022.

HILLESHEIM, J.C. *et al.* 2026 Overview of the physics basis for the ARC fusion power plant. *J. Plasma Phys.* **92**, E69. DOI: 10.1017/S0022377826101706

HO, A., CITRIN, J., AURIEMMA, F., BOURDELLE, C., CASSON, F., KIM, H.-T., MANAS, P., SZEPESI, G., WEISEN, H. & CONTRIBUTORS, J. 2019 Application of gaussian process regression to plasma turbulent transport model validation via integrated modelling. *Nucl. Fusion* **59**, 056007.

HOLLAND, C., BASS, E., ORLOV, D., MCCLENAGHAN, J., LYONS, B., GRIERSON, B., JIAN, X., HOWARD, N. & RODRIGUEZ-FERNANDEZ, P. 2023 Development of compact tokamak fusion reactor use cases to inform future transport studies. *J. Plasma Phys.* **89**, 905890418.

HOLLAND, C., WHITE, A.E., MCKEE, G.R., SHAFER, M.W., CANDY, J., WALTZ, R.E., SCHMITZ, L. & TYNAN, G.R. 2009 Implementation and application of two synthetic diagnostics for validating simulations of core tokamak turbulence. *Phys. Plasmas* **16**, 052301.

HOULBERG, W., ATTENBERGER, S. & HIVELY, L. 1982 Contour analysis of fusion reactor plasma performance. *Nucl. Fusion* **22**, 935.

HOWARD, N., GREENWALD, M., MIKKELSEN, D., REINKE, M., WHITE, A., ERNST, D., PODPALY, Y. & CANDY, J. 2012 Quantitative comparison of experimental impurity transport with nonlinear gyrokinetic simulation in an alcator c-mod l-mode plasma. *Nucl. Fusion* **52**, 063002.

HOWARD, N., RODRIGUEZ-FERNANDEZ, P., HOLLAND, C. & CANDY, J. 2024*a* Prediction of performance and turbulence in iter burning plasmas via nonlinear gyrokinetic profile prediction. *Nucl. Fusion* **65**, 016002.

HOWARD, N.T. *et al.* 2024*b* Simultaneous reproduction of experimental profiles, fluxes, transport coefficients, and tur*b*ulence characteristics via nonlinear gyrokinetic profile predictions in a DIII-d ITER similar shape plasma. *Phys. Plasmas* **31**, 032501.

HOWARD, N.T., HOLLAND, C., WHITE, A.E., GREENWALD, M., CANDY, J. & CREELY, A.J. 2016 Multi-scale gyrokinetic simulations: comparison with experiment and implications for predicting turbulence and transport. *Phys. Plasmas* **23**, 056109.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

HOWARD, N.T., RODRIGUEZ-FERNANDEZ, P., HOLLAND, C., RICE, J.E., GREENWALD, M., CANDY, J. & SCIORTINO, F. 2021 Gyrokinetic simulation of turbulence and transport in the SPARC tokamak. *Phys. Plasmas* **28**, 072502.

HU, Q.M., NAZIKIAN, R., GRIERSON, B.A., LOGAN, N.C., ORLOV, D.M., PAZ-SOLDAN, C. & YU, Q. 2020 Wide operational windows of edge-localized mode suppression by resonant magnetic perturbations in the diii-d tokamak. *Phys. Rev. Lett.* **125**, 045001.

HUGHES, J. *et al.* 2018 Access to pedestal pressure relevant to burning plasmas on the high magnetic field tokamak alcator c-mod. *Nucl. Fusion* **58**, 112003.

IVANOVA-STANIK, I., CHMIELEWSKI, P. & ZAGORSKI, R. 2023 Integrated core-sol simulations for sparc tokamak with the corediv code. *Fusion Engng Des.* **193**, 113698.

KINSEY, J., STAEBLER, G., CANDY, J., WALTZ, R. & BUDNY, R. 2011 ITER predictions using the GYRO verified and experimentally validated trapped gyro-landau fluid transport model. *Nucl. Fusion* **51**, 083001.

KOTSCHENREUTHER, M. *et al.* 2019 Gyrokinetic analysis and simulation of pedestals to identify the culprits for energy losses using 'fingerprints. *Nucl. Fusion* **59**, 096001.

MARTIN, Y.R., TAKIZUKA, T. & (and the ITPA CDBM H-Mode Threshold Database Working Group) 2008 Power requirement for accessing the h-mode in iter. *J. Phys.: Confer. Ser.* **123**, 012033.

MENEGHINI, O. *et al.* 2020 Neural-network accelerated coupled core-pedestal simulations with self-consistent transport of impurities and compatible with iter imas. *Nucl. Fusion* **61**, 026006.

MURACA, M., RODRIGUEZ-FERNANDEZ, P., HOWARD, N.T., HALL, J., FABLE, E. & TARDINI, G. 2025 Integrated modeling of sparc h-mode scenarios: exploration of the impact of modeling assumptions on predicted performance. arXiv: 2502.00187.

NAVARRO, A.B., HAPPEL, T., GÖRLER, T., JENKO, F., ABITEBOUL, J., BUSTOS, A., DOERK, H. & TOLD, D. 2015 Gyrokinetic studies of core turbulence features in asdex upgrade h-mode plasmas. *Phys. Plasmas* **22**, 042513.

NEU, R. *et al.* 2005 Tungsten: an option for divertor and main chamber plasma facing components in future fusion devices. *Nucl. Fusion* **45**, 209.

NEU, R. *et al.* 2016 Advanced tungsten materials for plasma-facing components of demo and fusion power plants. *Fusion Engng Des.* **109–111**, 1046–1052. Proceedings of the 12th International Symposium on Fusion Nuclear Technology-12 (ISFNT-12).

ODSTRČIL, T., HOWARD, N.T., SCIORTINO, F., CHRYSTAL, C., HOLLAND, C., HOLLMANN, E., MCKEE, G., THOME, K.E. & WILKS, T.M. 2020 Dependence of the impurity transport on the dominant turbulent regime in elm-y h-mode discharges on the DIII-D tokamak. *Phys. Plasmas* **27**, 082503.

PANKIN, A., MCCUNE, D., ANDRE, R., BATEMAN, G. & KRITZ, A. 2004 The tokamak monte carlo fast ion module nubeam in the national transport code collaboration library. *Comput. Phys. Commun.* **159**, 157–184.

PEREVERZEV, G., YUSHMANOV, P., D.A.P.A.T., K. & Z., L. 1991 Astra - an automatic system for transport analysis in a tokamak. *Tech. Rep*. Germany iPP–5/42 INIS Reference Number: 23024293.

PORCELLI, F., BOUCHER, D. & ROSENBLUTH, M.N. 1996 Model for the sawtooth period and amplitude. *Plasma Phys. Control. Fusion* **38**, 2163.

PUTTERICH, T., NEU, R., DUX, R., WHITEFORD, A., O'MULLANE, M., SUMMERS, H. & the ASDEX upgrade team 2010 Calculation and experimental test of the cooling factor of tungsten. *Nucl. Fusion* **50**, 025012.

RODRIGUEZ-FERNANDEZ, P., HOWARD, N. & CANDY, J. 2022 Nonlinear gyrokinetic predictions of sparc burning plasma profiles enabled by surrogate modeling. *Nucl. Fusion* **62**, 076036.

RODRIGUEZ-FERNANDEZ, P., HOWARD, N., SALTZMAN, A., KANTAMNENI, S., CANDY, J., HOLLAND, C., BALANDAT, M., AMENT, S. & WHITE, A. 2024*a* Enh*a*ncing predictive capabilities in fusion burning plasmas through surrogate-based optimization in core transport solvers. *Nucl. Fusion* **64**, 076034.

RODRIGUEZ-FERNANDEZ, P., HOWARD, N.T., GREENWALD, M.J., CREELY, A.J., HUGHES, J.W., WRIGHT, J.C., HOLLAND, C., LIN, Y. & SCIORTINO, F. 2020 Predictions of core plasma performance for the sparc tokamak. *J. Plasma Phys.* **86**, 865860503.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press

RODRIGUEZ-FERNANDEZ, P., HOWARD, N.T., SALTZMAN, A., SHOJI, L., BODY, T., BATTAGLIA, D.J., HUGHES, J.W., CANDY, J., STAEBLER, G.M. & CREELY, A.J. 2024*b* Core performance predictions in projected SPARC first-campaign plasmas with nonlinear CGYRO. *Phys. Plasmas* **31**, 062501.

ROSS, D.W. & DORLAND, W. 2002 Comparing simulation of plasma turbulence with experiment. II. Gyrokinetic simulations. *Phys. Plasmas* **9**, 5031–5035.

RUIZ RUIZ, J. *et al.* 2025 Measurement of zero-frequency fluctuations generated by coupling between alfvén modes in the jet tokamak. *Phys. Rev. Lett.* **134**, 095103.

SEEBER, B. 2023 Prospects of an alternative superconductor technology for fusion reactors. *Phys. Plasmas* **30**, 120604.

SNYDER, P. *et al.* 2009 Pedestal stability comparison and iter pedestal prediction. *Nucl. Fusion* **49**, 085035.

SNYDER, P. *et al.* 2019 High fusion performance in Super H-mode experiments on Alcator C-Mod and DIII-D. *Nucl. Fusion* **59**, 086017.

SNYDER, P., GROEBNER, R., HUGHES, J., OSBORNE, T., BEURSKENS, M., LEONARD, A., WILSON, H. & XU, X. 2011 A first-principles predictive model of the pedestal height and width: development, testing and iter optimization with the eped model. *Nucl. Fusion* **51**, 103016.

SNYDER, P., OSBORNE, T., BURRELL, K., GROEBNER, R., LEONARD, A., NAZIKIAN, R., ORLOV, D., SCHMITZ, O., WADE, M. & WILSON, H. 2012 The EPED pedestal model and edge localized mode-suppressed regimes: studies of quiescent H-mode and development of a model for edge localized mode suppression via resonant magnetic perturbations. *Phys. Plasmas* **19**, 056115.

SNYDER, P., SOLOMON, W., BURRELL, K., GAROFALO, A., GRIERSON, B., GROEBNER, R., LEONARD, A., NAZIKIAN, R., OSBORNE, T. & BELLI, E. 2015 Super H-mode: theoretical prediction and initial observations of a new high performance regime for tokamak operation. *Nucl. Fusion* **55**, 083026.

SORBOM, B. *et al.* 2015 Arc: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets. *Fusion Engng Des.* **100**, 378–405.

STAEBLER, G.M., CANDY, J., BELLI, E.A., KINSEY, J.E., BONANOMI, N. & PATEL, B. 2020 Geometry dependence of the fluctuation intensity in gyrokinetic turbulence. *Plasma Phys. Control. Fusion* **63**, 015013.

STAEBLER, G.M., CANDY, J., HOWARD, N.T. & HOLLAND, C. 2016 The role of zonal flows in the saturation of multi-scale gyrokinetic turbulence. *Phys. Plasmas* **23**, 062518.

STAEBLER, G.M., KINSEY, J.E. & WALTZ, R.E. 2007 A theory-based transport model with comprehensive physicsa. *Phys. Plasmas* **14**, 055909.

SUGAMA, H., WATANABE, T.-H. & NUNAMI, M. 2009 Linearized model collision operators for multiple ion species plasmas and gyrokinetic entropy balance equations. *Phys. Plasmas* **16**, 112503.

VAN DE PLASSCHE, K.L., CITRIN, J., BOURDELLE, C., CAMENEN, Y., CASSON, F.J., DAGNELIE, V.I., FELICI, F., HO, A., VAN MULDERS, S. & CONTRIBUTORS JET 2026 Fast modeling of turbulent transport in fusion plasmas using neural networks. *Phys. Plasmas* **27** (2), 022310.

VERDOOLAEGE, G., KAYE, S.M., ANGIONI, C., KARDAUN, O.J.W.F., MASLOV, M., ROMANELLI, M., RYTER, F., THOMSEN, K., ASDEX Upgrade Team, the, EUROfusion MST1 Team, the, Contributors, J.E.T. 2021 The updated itpa global h-mode confinement database: description and analysis. *Nucl. Fusion* **61**, 076006.

WHITE, A.E. *et al.* 2008 Measurements of core electron temperature and density fluctuations in diii-d and comparison to nonlinear gyrokinetic simulations. *Phys. Plasmas* **15**, 056116.

WHITE, A.E. *et al.* 2013 Multi-channel transport experiments at alcator c-mod and comparison with gyrokinetic simulations. *Phys. Plasmas* **20**, 056106.

https://doi.org/10.1017/S0022377826101573 Published online by Cambridge University Press